arXiv:cs/0702088v1 [cs.GT] 16 Feb 2007

Paths Beyond Local Search: A Nearly Tight Bound for Randomized Fixed-Point Computation*

Xi Chen[†]
Department of Computer Science
Tsinghua University

Shang-Hua Teng[‡]
Department of Computer Science
Boston University

Abstract

In 1983, Aldous proved that randomization can speedup local search. For example, it reduces the query complexity of local search over $[1:n]^d$ from $\Theta(n^{d-1})$ to $O(d^{1/2}n^{d/2})$. It remains open whether randomization helps fixed-point computation. Inspired by this open problem and recent advances on equilibrium computation, we have been fascinated by the following question:

Is a fixed-point or an equilibrium fundamentally harder to find than a local optimum?

In this paper, we give a nearly-tight bound of $(\Omega(n))^{d-1}$ on the randomized query complexity for computing a fixed point of a discrete Brouwer function over $[1:n]^d$. Since the randomized query complexity of global optimization over $[1:n]^d$ is $\Theta(n^d)$, the randomized query model over $[1:n]^d$ strictly separates these three important search problems:

Global optimization is harder than fixed-point computation, and
fixed-point computation is harder than local search.

Our result indeed demonstrates that randomization does not help much in fixed-point computation in the query model; the deterministic complexity of this problem is $\Theta(n^{d-1})$.

---

[*] This research is supported mostly by the NSF ITR grant CCR-0325630.

[†] Part of this work was done while visiting Computer Science Department at Boston University. In part supported by the Chinese National Key Foundation R & D Plan (2003CB317807, 2004CB318108), the National Natural Science Foundation of China Grant 60553001 and the National Basic Research Program of China Grant (2007CB807900, 2007CB807901).

[‡] Part of this work was done while visiting Tsinghua University and Microsoft Research Asia Lab.

Prologue

Scene 1: *On the first day of your new job, your boss teaches you the Simplex Algorithm with the Steepest-Edge Pivoting Rule. You quickly master the steps of the algorithm. So she gives you a large linear program that simulates a new business model.*

"I am going to a convention in Hawaii for ten days. Could you work on the program starting with this initial vector $\mathbf{x}_0$*?" she asks. "The solution will be a vector that you cannot improve upon. Email it to me when you are done"*

So she leaves for beautiful Hawaii and you begin your iterative path-following search. Ten days later, she comes back, relaxed, right as you finish computing $\mathbf{x}_{1000000}$*!*

"I haven't found the solution yet," you report, "but I have followed the steepest-edges a million steps and get $\mathbf{x}_{1000000}$*."*

She takes the objective vector $\mathbf{c}$ *and quickly computes* $\mathbf{c}^T\mathbf{x}_{1000000}/\mathbf{c}^T\mathbf{x}_0$*, and it is* 1.10.

"You find a vector that is 10 percent better than what we had initially," she says cheerfully. "Good job!"

The next day, you get a ten-percent raise.

Scene 2: *On the first day of your new job, your boss teaches you the Lemke-Howson algorithm for finding a Nash equilibrium of a two-player game. You quickly master the steps of the algorithm. So he gives you a large two-player game that models a two-group exchange market.*

"I am going to a convention in Hawaii for ten days. Could you work on this two-player game?" he asks. "Here is an initial strategy-profile," he gives you $(\mathbf{x}_0, \mathbf{y}_0)$*, "and the strategy-profile that Lemke-Howson halts on is a Nash equilibrium. Email it to me when you are done."*

So he leaves for beautiful Hawaii and you begin your iterative path-following search. Ten days later, he comes back, relaxed, right as you finish computing $(\mathbf{x}_{1000000}, \mathbf{y}_{1000000})$*!*

"I haven't found the solution yet," you report, "but I have followed the Lemke-Howson path a million steps and get $(\mathbf{x}_{1000000}, \mathbf{y}_{1000000})$*."*

He looks at $(\mathbf{x}_{1000000}, \mathbf{y}_{1000000})$ *for a while and then frowns, just slightly.*

"Hmmmm, no equilibrium in a million steps!" he says. "Well, good job and thanks."

The next day, you still have your job but get no raise.

1 Introduction

The Simplex Algorithm [11] is an example of an implementation of local search[1] and finding a Nash equilibrium [22] is an example of fixed-point computation (FPC). A general approach for local search is Iterative Improvement. Steepest-Descent is its most popular example. It follows a path in the feasible space, a path along which the objective values are monotonically improving. The end of the path is a local optimum. Like Iterative Improvement, many algorithms for FPC, such as the Lemke-Howson algorithm [20] and the constructive proof of Sperner's Lemma [29], also follow a path whose endpoint is an equilibrium or a fixed-point. But unlike a path in local search, a path in FPC does not have an obvious "locally computable" monotonic[2] measure-of-progress. Moreover, path following in FPC from an arbitrary point could lead to a cycle while the union of paths in Iterative Improvement is acyclic.

Do these structural differences have any algorithmic implication?

There have been increasing evidence, beyond the stories of our prologue, that local search and FPC are very different. First, Aldous [2] showed that randomization can speedup local search (more discussion below). His method crucially utilizes the monotonicity discussed above. It remains open whether randomization helps FPC. Second, polynomial-time path-following-like algorithms have been developed for some non-trivial classes of local search problems. These algorithms include the interior-point algorithm for linear and convex programming [18, 23] and edge-insertion algorithms for geometric optimization [13]. However, popular fixed-point problems, such as the computation of a Nash or a market equilibrium [3] might be hard for polynomial time [12, 7, 10]. Other than those that can be solved by convex programming, we haven't yet discovered a significantly non-trivial class of equilibrium problems that are solvable in polynomial-time. Third, an approximate local optimum for every PLS (Polynomial Local Search) problem can be found in fully-polynomial time [24]. In contrast, although a faster randomized algorithm was found for approximating Nash equilibria [21], finding an approximate Nash equilibrium in fully-polynomial time is computationally equivalent to finding an exact Nash equilibrium in polynomial time [8]. We face the same challenge in approximating market equilibria [16]. Fourth, although they all have exponential worst-case complexity [27, 19], the smoothed complexity of the Simplex Algorithm and Lemke-Howson Algorithm (or Scarf's market equilibrium algorithm [28]) might be drastically different [30, 8, 16]. This evidence inspires us to ask:

Is fixed-point computation fundamentally harder than local search?

To investigate this question, we consider the complexity of these two search problems defined over $\mathbb{Z}_n^d = [1:n]^d$. For fixed-points, we are given a function $F : \mathbb{Z}_n^d \rightarrow \mathbb{Z}_n^d$ that satisfies Brouwer's condition [4] — a set of continuity and boundary conditions (see Section 2) — that guarantees the existence of a fixed-point. Recall that a vector $\mathbf{v} \in \mathbb{Z}_n^d$ is a fixed-point of F if $F(\mathbf{v}) = \mathbf{v}$. The *FPC problem* is to find a fixed-point of F. For local optima, we are given a function $h : \mathbb{Z}_n^d \to \mathbb{R}$. The *local search problem* is to find a local optimum of h, for example, a vector $\mathbf{x} \in \mathbb{Z}_n^d$ such that $h(\mathbf{x}) \geq h(\mathbf{y})$, $\forall \mathbf{y}$ with $||\mathbf{x} - \mathbf{y}||_1 \leq 1$.

---

[1] Note that in linear programming, each local optimum is also a global optimum.

[2] Each path has a "globally computable" monotonic measure, the number of hops from the start of the path to a node.

For both problems, we consider the query complexity in the query model: The algorithm can only access F and h, respectively, by asking queries of the form: "What is $F(\mathbf{x})$?" and "What is $h(\mathbf{x})$?". The complexity is measured by the number of queries needed to find a solution.

There are some similarities between FPC and local search over $\mathbb{Z}_n^d$. For both, divide-and-conquer has positive but limited success: Both problems can be solved by $O(n^{d-1})$ queries [5]. An alternative approach to solve both problems is path-following. When following a short path, it can be faster than divide-and-conquer. But for both problems, long and winding paths are the cause of inefficiency.

However, there is one prominent difference between a path to a local optimum and a path to a fixed point. The values of h along a path to a local optimum are monotonic, serving as a measure-of-progress along the path. Aldous [2] used this fact in a randomized algorithm: Randomly query $d^{1/2}n^{d/2}$ points in $\mathbb{Z}_n^d$; let $\mathbf{s}$ be the sample point with the largest h value; follow a path starting at $\mathbf{s}$. If a path to a local optimum is long, say much longer than $d^{1/2}n^{d/2}$, then with high probability, the random samples intersect the path and partition it into sub-paths, each with expected length $O(d^{1/2}n^{d/2})$. As $\mathbf{s}$ has the largest h value, its sub-path is the last sub-path of a potentially long path, and we expect its length to be $O(d^{1/2}n^{d/2})$. So with randomization, Aldous reduced the expected query complexity to $O(d^{1/2}n^{d/2})$.

But it remains open whether randomization can reduce the query complexity of FPC over $\mathbb{Z}_n^d$. The lack of a measure-of-progress along a path makes it impossible for us to directly use Aldous' idea.

Our Main Result

The state of our knowledge suggests that FPC might be significantly harder than local search, at least in the randomized query model. We have formulated a concrete conjecture stating that an expected number of $(\Omega(n))^{d-1}$ queries are needed in randomized FPC over $\mathbb{Z}_n^d$.

As the main technical result of this paper, we prove that an expected number of $(\Omega(n))^{d-1}$ queries are indeed needed. Our lower bound is essentially tight[3], since the deterministic divide-and-conquer algorithm in [5] can find a fixed point by querying $O(n^{d-1})$ vectors. In contrast to Aldous's result [2], our result demonstrates that randomization does not help much in FPC in the query model. It shows that, in the randomized query model over $\mathbb{Z}_n^d$, a fixed-point is strictly harder to find than a local optimum! The significant gap between these two problems is revealed only in randomized computation. In the deterministic framework, both have query complexity $\Theta(n^{d-1})$.

One can show that the randomized query complexity for finding a global optimum over $\mathbb{Z}_n^d$ is $\Theta(n^d)$. So, the randomized query model over $\mathbb{Z}_n^d$ strictly separates these three important search problems:

Global optimization is harder than fixed-point computation, and
fixed-point computation is harder than local search.

We anticipate that a similar gap can be obtained in the quantum query model.

---

[3] The constant in Ω in our lower bound depends exponentially on d. See Theorem 2.2.

Related Work and Technical Contributions

Our work is also inspired by the lower bound results of Aaronson [1], Santha and Szegedy [26], Zhang [33], and Sun and Yao [31] on the randomized and quantum query complexity of local search over $\mathbb{Z}_n^d$.

In this paper, we introduce several new techniques to study the complexity of FPC. Instrumental to our analysis, we develop a method to generate hard-to-find random long paths in the grid graph over $\mathbb{Z}_n^d$. To achieve our nearly-tight lower bound, these paths must be much longer than the random paths constructed in [33, 31] for local search. Our paths has expected length $(\Theta(n))^{d-1}$ while those random paths for local search have length $\Theta(n^{d/2})$. We also develop new techniques for unknoting a self-intersecting path and for realizing a path with a Brouwer function. These techniques might be useful on their own in the future algorithmic and complexity-theoretic studies of FPC and its applications.

There are several earlier work on the query complexity of FPC. Hirsch, Papadimitriou and Vavasis [15] considered the deterministic query complexity of FPC. They proved a tight $\Theta(n)$ bound for $\mathbb{Z}_n^2$ and an $\Omega(n^{d-2})$ lower bound for $\mathbb{Z}_n^d$. Subsequently, Chen and Deng [5] improved this bound to $\Theta(n^{d-1})$ for $\mathbb{Z}_n^d$. Recently, Friedl, Ivanyos, Santha, and Verhoeven [14] gave a $\Omega(n^{1/4})$-lower bound on the randomized query-complexity of the 2D Sperner problem. Our method for unkonting self-intersecting paths can be viewed as an extension of the 2D technique of [6] to high dimensions.

Paper Organization

In Section 2, we introduce three high-dimensional search problems. In Section 3, we reduce one of them, called End-of-a-String, to fixed-point computation over $\mathbb{Z}_n^d$. In Section 4, we give a nearly tight bound on the randomized query complexity of End-of-a-String. Together with the reduction in Section 3, we obtain our main result on fixed-point computation.

2 Three High-Dimensional Search Problems

We will define three search problems. The first one concerns FPC. We introduce the last two to help the study of the first one. Below, let $\mathbb{E}^d = \{\pm\mathbf{e}_1, \pm\mathbf{e}_2, ..., \pm\mathbf{e}_d\}$ be the set of *principle unit-vectors* in d-dimensions. Let $\|\cdot\|$ denote $\|\cdot\|_\infty$. For two vectors[4] $\mathbf{u} \neq \mathbf{v}$ in $\mathbb{Z}^d$, we say $\mathbf{u} < \mathbf{v}$ lexicographically if $u_i < v_i$ and $u_j = v_j$ for all $1 \leq j < i$, for some i.

For each of the three search problems, we will define its mathematical structure, a query model for accessing this structure, the search problem itself, and its query complexity.

2.1 Discrete Brouwer Fixed-Points

Recall that a vector $\mathbf{v} \in \mathbb{Z}_n^d$ is a *fixed-point* of a function F from $\mathbb{Z}_n^d$ to $\mathbb{Z}_n^d$ if $F(\mathbf{v}) = \mathbf{v}$. A function $f : \mathbb{Z}_n^d \rightarrow \{\mathbf{0}\} \cup \mathbb{E}^d$ is *bounded* if $f(\mathbf{x}) + \mathbf{x} \in \mathbb{Z}_n^d$ for all $\mathbf{x} \in \mathbb{Z}_n^d$; $\mathbf{v} \in \mathbb{Z}_n^d$ is a *zero point* of f if $f(\mathbf{v}) = \mathbf{0}$. Clearly, if $F(\mathbf{x}) = \mathbf{x} + f(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{Z}_n^d$, then $\mathbf{v}$ is a fixed point of F iff $\mathbf{v}$ is a zero point of f.

---

[4] We will use bold lower-case Roman letters such as $\mathbf{x}$, $\mathbf{a}$, $\mathbf{b}_j$ to denote vectors. Whenever a vector, say $\mathbf{a} \in \mathbb{R}^n$ is present, its components will be denoted by lower-case Roman letters with subscripts, such as $a_1, \ldots, a_n$. So entries of $\mathbf{b}_j$ are $(b_{j,1}, \ldots, b_{j,n})$.

Definition 2.1 (Direction Preserving Functions)**.** *A function f from S to $\{\mathbf{0}\} \cup \mathbb{E}^d$ where $S \subset \mathbb{Z}^d$ is* direction-preserving *if $\|f(\mathbf{r}_1) - f(\mathbf{r}_2)\| \leq 1$ for all pairs $\mathbf{r}_1, \mathbf{r}_2 \in S$ such that $\|\mathbf{r}_1 - \mathbf{r}_2\| \leq 1$.*

Following the discrete fixed-point theorem of [17], we have: For every function $f : \mathbb{Z}_n^d \rightarrow \{\mathbf{0}\} \cup \mathbb{E}^d$, if f is both bounded and direction-preserving, then there exists $\mathbf{v} \in \mathbb{Z}_n^d$ such that $f(\mathbf{v}) = \mathbf{0}$. We refer to a bounded and direction-preserving function f over $\mathbb{Z}_n^d$ as a *Discrete Brouwer function* or simply a *Brouwer function* over $\mathbb{Z}_n^d$. In the query model, one can only access f by asking queries of the form: "What is $f(\mathbf{r})$?" for a query point $\mathbf{r} \in \mathbb{Z}_n^d$.

The FPC problem ZP^d that we will study is as follows: *Given a Brouwer function f from $\mathbb{Z}_n^d$ to $\{\mathbf{0}\} \cup \mathbb{E}^d$ in the query model, find a zero point of f.* Let $\mathsf{RQ}_{\mathsf{ZP}}(f)$ denote the expected number of queries needed by the best randomized algorithm to find[5] a zero point of f. We let

$$\mathsf{RQ}_{\mathsf{ZP}}^d(n) = \max_{f\text{: Brouwer function over } \mathbb{Z}_n^d} \big\{ \mathsf{RQ}_{\mathsf{ZP}}(f) \big\},$$

be the *randomized query complexity* for solving ZP^d. In this paper, we will prove:

Theorem 2.2 (Randomized Query Complexity of Fixed Points)**.** *There is a constant c such that for all sufficiently large n,*

$$\mathsf{RQ}_{\mathsf{ZP}}^d(n) \geq \left(\frac{n}{c^d}\right)^{d-1}.$$

In contrast, the deterministic query complexity for solving ZP^d is at most $7n^{d-1}$ [5]. The Brouwer fixed point problem defined here is computationally equivalent to the fixed problems defined in [15, 12, 8]. Thus, our result carries over to these FPC problems.

2.2 End-of-a-Path in Grid-PPAD Graphs

The mathematical structure for this search problem is a directed graph $G = (V, E)$. A vertex $v \in V$ satisfies *Euler's condition* if $\Delta_I(v) = \Delta_O(v)$ where $\Delta_I(v)$ and $\Delta_O(v)$ are the in-degree and the out-degree of v. We start with the following definition motivated by Papadimitriou's **PPAD** class [25].

Definition 2.3 (Generalized PPAD Graphs)**.** *A directed graph $G = (V, E)$ is a* generalized PPAD graph *if (1) there exists exactly one vertex $v_S \in V$ with $\Delta_O(v_S) = \Delta_I(v_S) + 1$ and exactly one vertex $v_T \in V$ with $\Delta_I(v_T) = \Delta_O(v_T) + 1$. (2) all vertices in $V - \{v_S, v_T\}$ satisfy Euler's condition and (3) if (v_1, v_2) is a directed edge in E, then $(v_2, v_1) \notin E$. We refer to v_S and v_T as the* starting *and* ending *vertices of G, respectively.*

We call G a PPAD graph *if in addition $\Delta_I(v), \Delta_O(v) \leq 1$, for all $v \in V$.*

Edges of a PPAD graph form a collection of disjoint directed cycles and a directed path from v_S to v_T. In this paper, we are interested in a special family of PPAD graphs over $\mathbb{Z}_n^d$. A directed graph $G = (\mathbb{Z}_n^d, E)$ is a generalized grid PPAD-graph over $\mathbb{Z}_n^d$ if it is a generalized

---

[5] One can also change "to find" to "to find, with high probability".

PPAD graph and the underlying undirected graph of G is a subgraph of the grid graph defined over $\mathbb{Z}_n^d$. Moreover, if G is also a PPAD graph, then we say G is a grid PPAD graph.

We now define the query model $\mathbb{B}_G$ for accessing a grid PPAD graph G.

Definition 2.4 ($\mathbb{B}_G$)**.** *$\mathbb{B}_G$ is a map from $\mathbb{Z}_n^d$ to $(\{\text{“no”}\} \cup \mathbb{E}^d) \times (\{\text{“no”}\} \cup \mathbb{E}^d)$ such that, for all $\mathbf{v} \in \mathbb{Z}_n^d$,*

- *$\mathbb{B}_G(\mathbf{v}) = (\text{“no”}, \mathbf{v}_1 - \mathbf{v})$ if $\mathbf{v}$ is the starting vertex of G and $(\mathbf{v}, \mathbf{v}_1) \in E$;*
- *$\mathbb{B}_G(\mathbf{v}) = (\mathbf{v} - \mathbf{v}_1, \text{“no”})$ if $\mathbf{v}$ is the ending vertex of G and $(\mathbf{v}_1, \mathbf{v}) \in E$;*
- *$\mathbb{B}_G(\mathbf{v}) = (\mathbf{v} - \mathbf{v}_1, \mathbf{v}_2 - \mathbf{v})$ if $(\mathbf{v}_1, \mathbf{v})$ and $(\mathbf{v}, \mathbf{v}_2)$ are directed edges of G.*
- *$\mathbb{B}_G(\mathbf{v}) = (\text{“no”}, \text{“no”})$, otherwise.*

In other words, $\mathbb{B}_G$ specifies the predecessor and successor of each vertex $\mathbf{v}$ in G. We will use the property that if $\mathbb{B}_G(\mathbf{v}) = (\mathbf{s}_1, \mathbf{s}_2)$ and $\mathbf{s}_1, \mathbf{s}_2 \in \mathbb{E}^d$, then $\mathbf{s}_1 + \mathbf{s}_2 \neq 0$.

Let GP^d be the search problem: *Given a triple $(G, 0^n, \mathbf{u})$, where G is a grid PPAD graph over $\mathbb{Z}_n^d$ accessible by $\mathbb{B}_G$ and $\mathbf{u}$ is the starting vertex of graph G satisfying $u_d = 1$, find its ending vertex.* We use $\mathsf{RQ}_{\mathsf{GP}}^d(n)$ to denote the randomized query complexity for solving this problem.

2.3 End-of-a-String

Suppose Σ is a finite set. A string S over Σ of length m is a sequence $S = a_1 a_2 ... a_{m-1} a_m$ with $a_i \in \Sigma$. We use $|S| = m$ to denote the length of S.

Definition 2.5 (Non-Repeating-Strings)**.** *A string $S = a_1 a_2 ... a_m$ over $\mathbb{Z}_n = [1:n]$ is d-*non-repeating *for $d \in [1:m]$, if (1) each string over $\mathbb{Z}_n$ of length d appears in S at most once; (2) a_i is odd if i is a multiple of d and a_i is even otherwise; and (3) m is a multiple of d. We define $\mathsf{end}_d(S) = a_{m-d+1} ... a_m$.*

Each d-non-repeating string $S = a_1 ... a_m$ over $\mathbb{Z}_n$ defines a query oracle $\mathbb{B}_S$ from $\mathbb{Z}_n^d$ to $(\{\text{“no”}\} \cup \mathbb{Z}_n) \times (\{\text{“no”}\} \cup \mathbb{Z}_n)$: For $S' = b_1 b_2 ... b_d \in \mathbb{Z}_n^d$, if S' is not a substring of S, then $\mathbb{B}_S(S') = (\text{“no”}, \text{“no”})$; otherwise, there is a unique k such that $a_{k+i-1} = b_i$, $\forall\, i \in [1:d]$. Then $\mathbb{B}_S(S') = (\text{“no”}, a_{d+1})$ if $k = 1$, $\mathbb{B}_S(S') = (a_{m-d}, \text{“no”})$ if $k = m - d + 1$, i.e., $S' = \mathsf{end}_d(S)$, and $\mathbb{B}_S(S') = (a_{k-1}, a_{k+d})$, otherwise.

Let ES^d be the search problem: *Given a d-non-repeating string S over $\mathbb{Z}_n$ accessible by $\mathbb{B}_S$, and its first d symbols $a_1 a_2 ... a_d$ where $a_d = 1$, find $\mathsf{end}_d(S)$.* We let $\mathsf{RQ}_{\mathsf{ES}}^d(n)$ denote its randomized query complexity. It is easy to show that $\mathsf{RQ}_{\mathsf{ES}}^1(n) = \Theta(n)$. In section 4, we will prove

Theorem 2.6 (Complexity of ES^d)**.** *For all sufficiently large n,*

$$\mathsf{RQ}_{\mathsf{ES}}^d(4n+4) \geq \frac{1}{2} \left(\frac{n}{2 \cdot 24^d} \right)^d.$$

3 Reduction Among Search Problems

In this section, we reduce ES^{d-1} to ZP^d by first reducing ES^{d-1} to GP^d (Theorem 3.1 below) and then reducing GP^d to ZP^d (Theorem 3.2). Theorem 2.2 then follows from Theorem 2.6.

Theorem 3.1 (From ES^{d-1} to GP^d)**.** *For all* $d \geq 2$, $\mathsf{RQ}_{\mathsf{ES}}^{d-1}(n) \leq 4d \cdot \mathsf{RQ}_{\mathsf{GP}}^{d}(8n+1)$.

Theorem 3.2 (From GP^d to ZP^d)**.** *For all* $d \geq 1$, $\mathsf{RQ}_{\mathsf{GP}}^{d}(n) \leq \mathsf{RQ}_{\mathsf{ZP}}^{d}(24n+7)$.

3.1 From ES^{d-1} to GP^d: Proof of Theorem 3.1

Proof. [of Theorem 3.1]: We define a map $\mathcal{F}_d$ from $\mathbb{Z}^{d-1}$ to $\mathbb{Z}^d$: for $d = 2$, $\mathcal{F}_2(a) = (a, a)$; and for $d > 2$, $\mathcal{F}_d(\mathbf{a}) = (a_1, a_1 + a_2, ..., a_{d-2} + a_{d-1}, a_{d-1})$. We will crucially use the following nice property of $\mathcal{F}_d$.

> For any $k \in [1 : d]$ and for any $\mathbf{a} \in \mathbb{Z}^{d-1}$, we can uniquely determine the first k and the last k entries of $\mathbf{a}$, respectively, from the first k and the last k entries of $\mathcal{F}_d(\mathbf{a})$.

Let S be a $(d-1)$-non-repeating string over $\mathbb{Z}_n$ of length $m(d-1)$ for some $m \geq 2$, whose $(d-1)^{st}$ symbol is 1. We view S as a sequence of m points $\mathbf{a}_1, \mathbf{a}_2, ...\mathbf{a}_m$ in $\mathbb{Z}_n^{d-1}$, where $\mathbf{a}_i = a_{i,1}...a_{i,d-1}$, such that, $S = a_{1,1}a_{1,2}...a_{1,d-1}...a_{m,1}a_{m,2}...a_{m,d-1}$. From S, we will construct a grid PPAD graph G' in two stages. In the first stage, we construct a generalized grid PPAD graph G^* over $\mathbb{Z}_{2n}^d$ such that

- **(A.1)** Its starting vertex is $\mathbf{u}^* = \mathcal{F}_d(\mathbf{a}_1)$ and its ending vertex is $\mathbf{w}^* = \mathcal{F}_d(\mathbf{a}_m)$;
- **(A.2)** For every directed edge $(\mathbf{u}, \mathbf{v})$ with $\mathbf{u} - \mathbf{v} \in \mathbb{E}^d$, at most one query to $\mathbb{B}_S$ is needed to determine whether $(\mathbf{u}, \mathbf{v}) \in G^*$.

Recall that a directed path is *simple* if it contains each vertex at most once. Suppose $\mathbf{u}, \mathbf{v} \in \mathbb{Z}_{2n}^d$ are two vertices that differ in only one coordinate, say the i^{th} coordinate. Suppose $\mathbf{e} = (\mathbf{v} - \mathbf{u})/|v_i - u_i| \in \mathbb{E}^d$. Let $E(\mathbf{u}, \mathbf{v}) = \{(\mathbf{u}, \mathbf{u} + \mathbf{e}), (\mathbf{u} + \mathbf{e}, \mathbf{u} + 2\mathbf{e}), ..., (\mathbf{v} - \mathbf{e}, \mathbf{v})\}$. For $n, m_1, m_2 \in \mathbb{Z}$ and $s \in \{\pm 1\}$, (n, s) is *consistent* with (m_1, m_2) if either $m_1 \leq n < m_2$ and $s = +1$ or $m_2 < n \leq m_1$ and $s = -1$.

We consider two consecutive points $\mathbf{a} = \mathbf{a}_t$ and $\mathbf{b} = \mathbf{a}_{t+1}$ in the $(d-1)$-non-repeating string S. We know $\mathbf{a} \neq \mathbf{b}$. We map them to vertices $\mathbf{u} = \mathcal{F}_d(\mathbf{a})$ and $\mathbf{w} = \mathcal{F}_d(\mathbf{b})$ in $\mathbb{Z}_{2n}^d$ and connect them with a path through a sequence of $(d-1)$ vertices $\mathbf{v}_0 = \mathbf{u}, \mathbf{v}_1, ..., \mathbf{v}_{d-1}, \mathbf{v}_d = \mathbf{w}$ where $v_{i,j} = u_j$ if $i < j$ and $v_{i,j} = w_j$ if $i \geq j$. Note that $\mathbf{v}_{i-1}$ and $\mathbf{v}_i$ differ only in the i^{th} coordinate. Let $P(\mathbf{a}, \mathbf{b}) = \cup_{i=0}^{d-1} E(\mathbf{v}_i, \mathbf{v}_{i+1})$. Then $P(\mathbf{a}, \mathbf{b})$ is a simple directed path in the grid graph over $\mathbb{Z}_{2n}^d$ from $\mathbf{u} = \mathbf{v}_0$ to $\mathbf{w} = \mathbf{v}_d$. As S is $(d-1)$-non-repeating, $\mathbf{a}_1 \neq \mathbf{a}_m$. By Property 3.3, Proposition 3.4 and Lemma 3.5 below, $G^* = (\mathbb{Z}_{2n}^d, \cup_{i=1}^{m-1} P(\mathbf{a}_i, \mathbf{a}_{i+1}))$ is a generalized grid PPAD graph. See Figure 1 for an example.

Proposition 3.3 (Path Union)**.** *Let* $P_1, P_2, ..., P_m$ *be* m *simple directed paths over* V *such that (1) each path has length at least one, (2) the ending vertex of* P_i *is same as the starting vertex of* P_{i+1}*, (3) the starting vertex of* P_1 *is different from the ending vertex of* P_m*, and (4) if* $(u, v) \in P_i$*, then* $(u, v), (v, u) \notin P_j$, $\forall j \neq i$*. Then* $G = (V, \cup_{i=1}^{m} P_i)$ *is a generalized PPAD graph.*

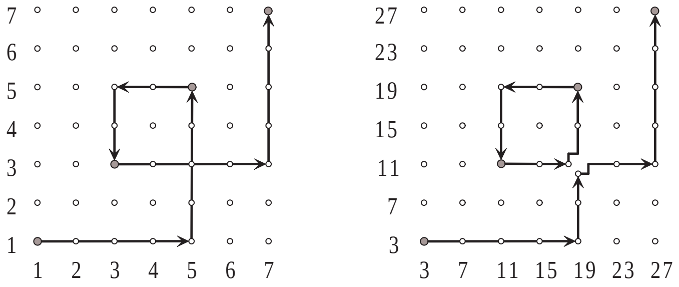

Figure 1: Graph G^* and G' constructed from string 1537

Proposition 3.4 (Local Characterization of $P(\mathbf{a}, \mathbf{b})$)**.** *For* $\mathbf{v} \in \mathbb{Z}_{2n}^d$ *and* $s \in \{\pm 1\}$*,*

1. $(\mathbf{v}, \mathbf{v} + s\mathbf{e}_1) \in P(\mathbf{a}, \mathbf{b})$ *if and only if* (v_1, s) *is consistent with* (a_1, b_1)*,* $a_{d-1} = v_d$*, and* $a_{d-i} = v_{d-i+1} - a_{d-i+1}$ *for all* $2 \leq i \leq d-1$*;*
2. $(\mathbf{v}, \mathbf{v} + s\mathbf{e}_d) \in P(\mathbf{a}, \mathbf{b})$ *if and only if* (v_d, s) *is consistent with* (a_{d-1}, b_{d-1})*,* $b_1 = v_1$*, and* $b_i = v_i - b_{i-1}$*, for* $2 \leq i \leq d-1$*; and*
3. *for* $1 < k < d$*,* $(\mathbf{v}, \mathbf{v} + s\mathbf{e}_k) \in P(\mathbf{a}, \mathbf{b})$ *if and only if* (v_k, s) *is consistent with* $(a_{k-1} + a_k, b_{k-1} + b_k)$ *and* (3.1) $a_{d-1} = v_d$*, and* $a_{d-i} = v_{d-i+1} - a_{d-i+1}$ *for* $2 \leq i \leq d-k$ *and* (3.2) $b_1 = v_1$*, and* $b_i = v_i - b_{i-1}$ *for* $2 \leq i \leq k-1$*.*

Lemma 3.5 (Structural Correctness)**.** *For all* $(d-1)$*-non-repeating string* $S = \mathbf{a}_1\mathbf{a}_2...\mathbf{a}_m$ *over* $\mathbb{Z}_n$*, if* $(\mathbf{u}, \mathbf{v}) \in P(\mathbf{a}_i, \mathbf{a}_{i+1})$ *then* $(\mathbf{u}, \mathbf{v}), (\mathbf{v}, \mathbf{u}) \notin P(\mathbf{a}_j, \mathbf{a}_{j+1})$ *for all* $i \neq j$*.*

Proof. We only prove the case when $\mathbf{e} = \mathbf{v} - \mathbf{u} = s\mathbf{e}_k$ with $1 < k < d$ and $s \in \{\pm 1\}$. The other two cases are similar. From Proposition 3.4, $(\mathbf{u}, \mathbf{v}) \in P(\mathbf{a}_i, \mathbf{a}_{i+1})$ implies that $\mathbf{a}_i$ and $\mathbf{a}_{i+1}$ satisfy conditions (3.1) and (3.2). If $(\mathbf{u}, \mathbf{v})$ or $(\mathbf{v}, \mathbf{u})$ is in $P(\mathbf{a}_j, \mathbf{a}_{j+1})$, then $\mathbf{a}_j$ and $\mathbf{a}_{j+1}$ also satisfy these two conditions. Then $a_{i,k}a_{i,k+1}...a_{i,d-1}a_{i+1,1}...a_{i+1,k-1} = a_{j,k}a_{j,k+1}...a_{j,d-1}a_{j+1,1}...a_{j+1,k-1}$, which contradicts with the assumption that S is $(d-1)$-non-repeating. ☐

We prove Property **A.2** as follows.

Proof of Property **A.2**. We will only prove for the case when $\mathbf{e} = \mathbf{v} - \mathbf{u} = s\mathbf{e}_k$ with $1 < k < d$. The other two cases are similar and simpler. To determine whether $(\mathbf{u}, \mathbf{v}) \in G^*$ or not, we consider the string $S' = a_k a_{k+1}...a_{d-1}b_1...b_{k-1}$ that satisfies both (3.1) and (3.2) in Proposition 3.4. Edge $(\mathbf{u}, \mathbf{v}) \in G^*$ if and only if 1) $S' \in \mathbb{Z}_n^{d-1}$; 2) a_{d-1} is odd; 3) $\mathbb{B}_S(S') = (a, b)$ for some $a, b \in \mathbb{Z}_n$; and 4) (u_k, s) is consistent with $(a + a_k, b_{k-1} + b)$, So, only one query to $\mathbb{B}_S$ is needed. ☐

In the second stage, we construct a grid PPAD graph G' over $\mathbb{Z}_{8n+1}^d$ from graph G^*. Let $\Gamma(\mathbf{v}) = 4\mathbf{v} - \mathbf{1}$ for all $\mathbf{v} \in \mathbb{Z}_{2n}^d$. Our G' will satisfy the following two properties. See Figure 1 for an example.

- **(B.1)** Its starting vertex is $\mathbf{u}' = \Gamma(\mathbf{u}^*) - 2\mathbf{e}_d$; its ending vertex $\mathbf{w}'$ satisfies $\|\mathbf{w}' - \Gamma(\mathbf{w}^*)\| \leq 1$;
- **(B.2)** For each $\mathbf{v} \in \mathbb{Z}_{8n+1}^d$, one can determine $\mathbb{B}_{G'}(\mathbf{v})$ from the predecessors and successors of $\mathbf{u}$ in G^*, where $\mathbf{u}$ is the lexicographically smallest vertex such that $\|\mathbf{v} - \Gamma(\mathbf{u}) \leq 2\|$.

Graph $G[H_1, H_2]$, where (H_1, H_2) is a balanced-non-canceling pair

```
set edge set $E[H_1, H_2] = \emptyset$
while $H_1 \neq \emptyset$ do
    let $\mathbf{s}_1$ be the smallest vector in $H_1$ and $\mathbf{s}_2$ be the largest vector in $H_2$
    according to the lexicographical ordering;
    set $H_1 = H_1 - \{\mathbf{s}_1\}$ and $H_2 = H_2 - \{\mathbf{s}_2\}$;
    set $E[H_1, H_2] = E[H_1, H_2] \cup \{(\mathbf{0} - \mathbf{s}_1, \mathbf{0} - \mathbf{s}_1 + \mathbf{s}_2), (\mathbf{0} - \mathbf{s}_1 + \mathbf{s}_2, \mathbf{0} + \mathbf{s}_2)\}$
```

Figure 2: Construction of Graph $G[H_1, H_2] = (\{-1, 0, +1\}^d, E[H_1, H_2])$

Two subsets H_1 and H_2 of $\mathbb{E}^d$, where $d \geq 2$, form a *balanced-non-canceling pair* if $|H_1| = |H_2|$ and $\mathbf{s}_1 + \mathbf{s}_2 \neq \mathbf{0}$ for all $\mathbf{s}_1 \in H_1$ and $\mathbf{s}_2 \in H_2$. Let $H_I(\mathbf{u}) = \{\mathbf{e} \in \mathbb{E}^d \mid (\mathbf{u} - \mathbf{e}, \mathbf{u}) \in E^*\}$ be the vector differences of $\mathbf{u}$ and its predecessors in G^*. Similarly, let $H_O(\mathbf{u}) = \{\mathbf{e} \in \mathbb{E}^d \mid (\mathbf{u}, \mathbf{u} + \mathbf{e}) \in E^*\}$ be the vector differences of the successors of $\mathbf{u}$ and $\mathbf{u}$. In the construction below, we will use the fact that if $\mathbf{u}$ satisfies Euler's condition then (H_I, H_O) is a balanced-non-canceling pair.

Using the procedure of Figure 2, we build a graph $G[H_1, H_2] = (\{-1, 0, +1\}^d, E[H_1, H_2])$ for each balanced-non-canceling pair H_1 and H_2. $G[H_1, H_2]$ has the following properties: (1) For every $\mathbf{u} \in \{-1, 0, +1\}^d$, $\Delta_I(\mathbf{u}), \Delta_O(\mathbf{u}) \leq 1$; (2) A vector $\mathbf{u} \in \{-1, 0, +1\}^d$ has $\Delta_I(\mathbf{u}) = 0$ and $\Delta_O(\mathbf{u}) = 1$ iff there exists an $\mathbf{e} \in H_1$ such that $\mathbf{u} = \mathbf{0} - \mathbf{e}$; (3) A vector $\mathbf{u} \in \{-1, 0, +1\}^d$ has $\Delta_I(\mathbf{u}) = 1$ and $\Delta_O(\mathbf{u}) = 0$ iff there exists an $\mathbf{e} \in H_2$ such that $\mathbf{u} = \mathbf{0} + \mathbf{e}$.

Let $\mathbf{u}^*$ be the starting vertex and $\mathbf{w}^*$ be the ending vertex of G^*. We build a grid PPAD graph $G' = (\mathbb{Z}^d_{8n+1}, E')$ by applying the procedure of Fig. 2 locally to every vertex $\mathbf{u} \in \mathbb{Z}^d_{2n}$ of G^*. We use $(H_I(\mathbf{u}), H_O(\mathbf{u}))$ or a slight modification of $(H_I(\mathbf{u}), H_O(\mathbf{u}))$ when $\mathbf{u} = \mathbf{u}^*$ or $\mathbf{w}^*$. Initially we set $E' = \emptyset$. Recall $\Gamma(\mathbf{u}) = 4\mathbf{u} - \mathbf{1}$.

1. [**local embedding of the starting vertex**] Since $u^*_d = 1$, we have $\mathbf{e}_d \notin H_I(\mathbf{u}^*)$ and $-\mathbf{e}_d \notin H_O(\mathbf{u}^*)$. Let $H_I = H_I(\mathbf{u}^*) \cup \{\mathbf{e}_d\}$. We add edges $(\Gamma(\mathbf{u}^*) - 2\mathbf{e}_d, \Gamma(\mathbf{u}^*) - \mathbf{e}_d)$ and $(\Gamma(\mathbf{u}^*) + \mathbf{s}_1, \Gamma(\mathbf{u}^*) + \mathbf{s}_2)$ to E' for all edges $(\mathbf{s}_1, \mathbf{s}_2)$ in $G[H_I, H_O(\mathbf{u}^*)]$.
2. [**local embedding of the ending vertex**] As $|H_I(\mathbf{w}^*)| = |H_O(\mathbf{w}^*)| + 1$, $H_I(\mathbf{w}^*) \neq \emptyset$. Let $\mathbf{e}$ be the smallest vector in $H_I(\mathbf{w}^*)$, and $H_I = H_I(\mathbf{w}^*) - \{\mathbf{e}\}$. Add edges $(\Gamma(\mathbf{w}^*) + \mathbf{s}_1, \Gamma(\mathbf{w}^*) + \mathbf{s}_2)$ to E' for all edges $(\mathbf{s}_1, \mathbf{s}_2)$ in $G[H_I, H_O(\mathbf{w}^*)]$.
3. [**local embedding of other vertices**] For each $\mathbf{u} \in G^*$, add $(\Gamma(\mathbf{u}) + \mathbf{s}_1, \Gamma(\mathbf{u}) + \mathbf{s}_2)$ to E' for all edges $(\mathbf{s}_1, \mathbf{s}_2)$ in $G[H_I(\mathbf{u}), H_O(\mathbf{u})]$.
4. [**connecting local embeddings**] For each edge $(\mathbf{u}, \mathbf{v}) \in G^*$, let $\mathbf{e} = \mathbf{v} - \mathbf{u} \in \mathbb{E}^d$. We add $(\Gamma(\mathbf{u}) + \mathbf{e}, \Gamma(\mathbf{u}) + 2\mathbf{e})$ and $(\Gamma(\mathbf{u}) + 2\mathbf{e}, \Gamma(\mathbf{u}) + 3\mathbf{e})$ to E'.

It is quite mechanical to check that G' is a PPAD grid graph that satisfies both Property **B.1** and **B.2**. We therefore complete the proof of Theorem 3.1. □

3.2 Canonicalization of Grid-PPAD Graphs

To ease our reduction from a grid-PPAD graph to a Brouwer function, we first canonicalize the grid-PPAD graph by regulating the way its path starts, moves, and ends.

Definition 3.6 (Canonical Grid-PPAD Graphs). *A grid-PPAD graph G over $\mathbb{Z}_n^d$ for $d \geq 2$ and $n > 1$ is* canonical *if $\mathbb{B}_G$ satisfies $\mathbb{B}_G(\mathbf{u}) \in \mathcal{S}^d$ for all $\mathbf{u} \in \mathbb{Z}_n^d$, where*

$$\mathcal{S}^2 = \{(\text{“}no\text{”}, \text{“}no\text{”}), (\text{“}no\text{”}, \mathbf{e}_2), (\mathbf{e}_2, \text{“}no\text{”})\} \cup \left\{ (\mathbf{s}_1, \mathbf{s}_2) \,\middle|\, \mathbf{s}_1, \mathbf{s}_2 \in \mathbb{E}^2, \mathbf{s}_1 + \mathbf{s}_2 \neq 0 \right\}.$$

For $d \geq 3$, $\mathcal{S}^d$ is the smallest subset of $(\{\text{“}no''\} \cup \mathbb{E}^d) \times (\{\text{“}no''\} \cup \mathbb{E}^d)$ that satisfies:

1. $(\text{“}no\text{”}, \text{“}no\text{”}), (\text{“}no\text{”}, \mathbf{e}_d), (\mathbf{e}_d, \text{“}no\text{”}) \in \mathcal{S}^d$*;*
2. $\{\mathbf{e}_d\} \times \{\mathbf{e}_{d-1}, \mathbf{e}_d\} \subset \mathcal{S}^d$ *and* $\{-\mathbf{e}_d\} \times \{\mathbf{e}_{d-1}, -\mathbf{e}_d\} \subset \mathcal{S}^d$*;*
3. $\{\mathbf{e}_k\} \times \{\mathbf{e}_{k-1}, \mathbf{e}_k, \pm\mathbf{e}_{k+1}\} \subset \mathcal{S}^d$ *and* $\{-\mathbf{e}_k\} \times \{\mathbf{e}_{k-1}, -\mathbf{e}_k\} \subset \mathcal{S}^d$*, for* $3 \leq k < d$*;*
4. $\{\mathbf{e}_2\} \times \{\pm\mathbf{e}_1, \mathbf{e}_2, \pm\mathbf{e}_3\} \subset \mathcal{S}^d$ *and* $\{-\mathbf{e}_2\} \times \{\pm\mathbf{e}_1, -\mathbf{e}_2\} \subset \mathcal{S}^d$*;*
5. $\{\mathbf{e}_1\} \times \{\mathbf{e}_1, \pm\mathbf{e}_2\} \subset \mathcal{S}^d$*; and* $\{-\mathbf{e}_1\} \times \{-\mathbf{e}_1, \pm\mathbf{e}_2\} \subset \mathcal{S}^d$*.*

Informally, edges in a canonical grid-PPAD graph over $\mathbb{Z}_n^d$ contains a single directed path starting at a point $\mathbf{u} \in \mathbb{Z}_n^d$ with $u_d = 1$ and ending at a point, say $\mathbf{w}$, and possibly some cycles. The second vertex on the path is $\mathbf{u} + \mathbf{e}_d$ and the second-to-the-last vertex is $\mathbf{w} - \mathbf{e}_d$. The path and the cycles satisfy the following conditions (below we will abuse “path” for both “path” and “cycle”): (1) To follow a directed edge along $\mathbf{e}_k$ (for $k \geq 3$), the path can only move locally in a 3D framework defined by $\{\mathbf{e}_{k-1}, \mathbf{e}_k, \pm\mathbf{e}_{k+1}\}$, see Figure 3, (for $k = d$, it can only move in a 2D framework). In a way, we view the d-dimensional space as a nested “affine subspaces” defined by $\{\pm\mathbf{e}_1, ..., \pm\mathbf{e}_k\}$ for $1 \leq k \leq d$. So to follow a positive principle direction $\mathbf{e}_k$, the path can *move down* a dimension along the positive direction $\mathbf{e}_{k-1}$, stay continuously along $\mathbf{e}_k$, or *move up* a dimension (unless $k = d$) along either $\pm\mathbf{e}_{k+1}$. (2) To follow a directed edge along $-\mathbf{e}_k$ for $k \geq 3$, the path can only move locally in a 2D framework defined by $\{\mathbf{e}_{k-1}, -\mathbf{e}_k\}$, see Figure 3. The path can *move down* a dimension along the positive direction of $\mathbf{e}_{k-1}$ or stay continuously along $-\mathbf{e}_k$, but it is not allowed to *move up* or *leave* this k-dimensional “affine subspace”. In the $\{\pm\mathbf{e}_1, \pm\mathbf{e}_2\}$ framework, the path is less restrictive as defined by conditions 4 and 5.

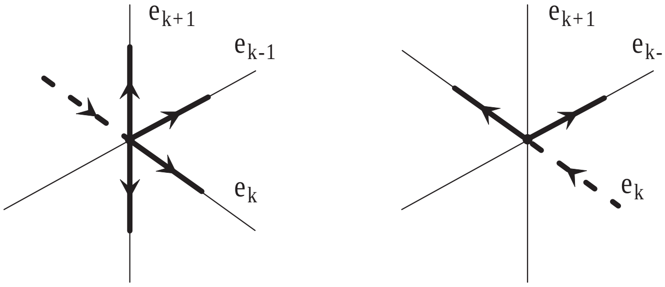

Figure 3:

In other words, the path can not move-up from an “affine subspace” (with the exception of the $\pm\mathbf{e}_1$ space) without first taking a step along the highest positive principle direction in the subspace. Similarly, the path can only move-down to an "affine subspace" by taking a positive first step along its highest principle direction. Otherwise, the path moves continuously.

Let CGP^d be: *Given a triple* $(G, 0^n, \mathbf{u}^*)$ *where* G *is a canonical grid-PPAD graph over* $\mathbb{Z}_n^d$ *accessible by* $\mathbb{B}_G$ *and* $\mathbf{u}^*$ *is the starting vertex of* G *with* $u_d^* = 1$*, find the ending vertex of* G*.* We use $\mathsf{RQ}^d_{\mathsf{CGP}}(n)$ to denote the randomized query complexity for solving this problem.

We now reduce GP^d to CGP^d. Before stating and proving this result, we first give two geometric lemmas. They provide the local operations for canonicalization.

We start with some notation. A sequence $P = \mathbf{u}_1 ... \mathbf{u}_m$, for $m \geq 2$, is a *canonical local path* if $\mathbf{u}_i$'s are distinct elements from $\{-3, -2, ..., +2, +3\}^d$ and $(\mathbf{u}_i - \mathbf{u}_{i-1}, \mathbf{u}_{i+1} - \mathbf{u}_i) \in \mathcal{S}^d$ for all $2 \leq i \leq m-1$. Suppose $P = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_m$ and $Q = \mathbf{v}_1\mathbf{v}_2 ... \mathbf{v}_k$ are two paths with $\mathbf{u}_m = \mathbf{v}_1$. We use $P \diamond Q$ to denote their "concatenation": $P \diamond Q = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_m\mathbf{v}_2 ... \mathbf{v}_k$.

Lemma 3.7 (Ending Gracefully)**.** *For each* $\mathbf{s} \in \mathbb{E}^d$ *with* $d \geq 2$*, there is a canonical local path* $P[d, \mathbf{s}] = \mathbf{u}_1 ... \mathbf{u}_m$ *satisfying* $\mathbf{u}_m - \mathbf{u}_{m-1} = \mathbf{e}_d$*,* $\mathbf{u}_1 = -3\mathbf{s}$*,* $\mathbf{u}_2 = -2\mathbf{s}$*, and* $\forall\, i \in [2:m]$*,* $\|\mathbf{u}_i\| \leq 2$*.*

Proof. We consider the three cases: $\mathbf{s} = \mathbf{e}_l$ for $1 \leq l \leq d$, $\mathbf{s} = -\mathbf{e}_1$ and $\mathbf{s} = -\mathbf{e}_l$ for $2 \leq l \leq d$.

In the first case, we set $P[d, \mathbf{s}] = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_{d-l+2}$ where $\mathbf{u}_1 = -3\mathbf{e}_l$, $\mathbf{u}_2 = -2\mathbf{e}_l$ and $\mathbf{u}_i = \mathbf{u}_{i-1} + \mathbf{e}_{l+i-2}$ for $3 \leq i \leq d-l+2$. In the second case, we set $P[d, \mathbf{s}] = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_d\mathbf{u}_{d+1}$ where $\mathbf{u}_1 = 3\mathbf{e}_1$, $\mathbf{u}_2 = 2\mathbf{e}_1$ and $\mathbf{u}_i = \mathbf{u}_{i-1} + \mathbf{e}_{i-1}$ for $3 \leq i \leq d+1$. In the third case, we set $P[d, \mathbf{s}] = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_{d-l+5}$ where $\mathbf{u}_1 = 3\mathbf{e}_l$, $\mathbf{u}_2 = 2\mathbf{e}_l$, $\mathbf{u}_3 = \mathbf{e}_l$ and $\mathbf{u}_i = \mathbf{u}_{i-1} + \mathbf{e}_{l+i-5}$ for $4 \leq i \leq d-l+5$.

One can easily check that $P[d, \mathbf{s}]$ satisfies the conditions of the lemma. □

Lemma 3.8 (Moving Gracefully)**.** *For all* $\mathbf{s}_1, \mathbf{s}_2 \in \mathbb{E}^d$ *with* $d \geq 2$ *such that* $\mathbf{s}_1 + \mathbf{s}_2 \neq 0$*, there exists a canonical local path* $P[d, \mathbf{s}_1, \mathbf{s}_2] = \mathbf{u}_1\mathbf{u}_2 ... \mathbf{u}_{m-1}\mathbf{u}_m$ *that satisfies* $\mathbf{u}_1 = -3\mathbf{s}_1$*,* $\mathbf{u}_2 = -2\mathbf{s}_1$*,* $\mathbf{u}_{m-1} = 2\mathbf{s}_2$*,* $\mathbf{u}_m = 3\mathbf{s}_2$*, and* $\forall\, i \in [2:m-1]$*,* $\|\mathbf{u}_i\| \leq 2$*.*

Proof. We prove by induction on d that there is a canonical local path $P[d, \mathbf{s}_1, \mathbf{s}_2]$ such that

1. $P[d, \mathbf{s}_1, \mathbf{s}_2]$ satisfies all the conditions in the statement of the lemma;
2. For all $\mathbf{s}_1, \mathbf{s}_2 \in \mathbb{E}^d$ such that $\mathbf{s}_1 + \mathbf{s}_2 \neq 0$, $\mathbf{e}_d, -\mathbf{e}_d \notin P[d, \mathbf{s}_1, \mathbf{s}_2]$; and
3. If $\mathbf{s}_1 = -\mathbf{e}_d$, then the first 3 vertices of $P[d, -\mathbf{e}_d, \mathbf{s}_2]$ are $3\mathbf{e}_d$, $2\mathbf{e}_d$ and $(\mathbf{e}_{d-1} + 2\mathbf{e}_d)$. If $\mathbf{s}_2 = -\mathbf{e}_d$, then the last 3 vertices of $P[d, \mathbf{s}_1, -\mathbf{e}_d]$ are $(-\mathbf{e}_{d-1} - 2\mathbf{e}_d)$, $-2\mathbf{e}_d$ and $-3\mathbf{e}_d$.

The base case when $d = 2$ is trivial. Inductively we assume, for $2 \leq d' < d$, path $P[d', \mathbf{s}'_1, \mathbf{s}'_2]$ exists for all $\mathbf{s}'_1, \mathbf{s}'_2 \in \mathbb{E}^{d'}$ such that $\mathbf{s}'_1 + \mathbf{s}'_2 \neq 0$. We let $P'[d', \mathbf{s}'_1, \mathbf{s}'_2]$ denote the sub-path of $P[d', \mathbf{s}'_1, \mathbf{s}'_2]$ such that

$$P[d', \mathbf{s}'_1, \mathbf{s}'_2] = (-3\mathbf{s}'_1)(-2\mathbf{s}'_1) \diamond P'[d', \mathbf{s}'_1, \mathbf{s}'_2] \diamond (2\mathbf{s}'_2)(3\mathbf{s}'_2).$$

Note that $P'[d', \mathbf{s}'_1, \mathbf{s}'_2]$ starts with $-2\mathbf{s}'_1$ and ends with $+2\mathbf{s}'_2$.

We will use D to denote the map $D(\mathbf{r}) = (r_1, ..., r_{d-1})$ from $\mathbb{Z}^d$ to $\mathbb{Z}^{d-1}$ and U to denote the map $U(\mathbf{r}) = (r_1, ..., r_{d-1}, 0)$ from $\mathbb{Z}^{d-1}$ to $\mathbb{Z}^d$. For a canonical local path $P = \mathbf{u}_1 ... \mathbf{u}_m$ in $\{-3, ..., +3\}^{d-1}$, we use $U(P)$ to denote path $U(\mathbf{u}_1)U(\mathbf{u}_2) ... U(\mathbf{u}_m)$ in $\{-3, ..., +3\}^d$.

For $\mathbf{s}_1, \mathbf{s}_2 \in \mathbb{E}^d$ with $\mathbf{s}_1 + \mathbf{s}_2 \neq 0$, we use the following procedure to build $P[d, \mathbf{s}_1, \mathbf{s}_2]$. Let

$$\begin{aligned} P^+ &= (-3\mathbf{e}_d)\,(-2\mathbf{e}_d)\,(\mathbf{e}_{d-1} - 2\mathbf{e}_d)\,(\mathbf{e}_{d-1} - \mathbf{e}_d)\,(\mathbf{e}_{d-1})\,(2\mathbf{e}_{d-1}); \\ P^- &= (3\mathbf{e}_d)\,(2\mathbf{e}_d)\,(\mathbf{e}_{d-1} + 2\mathbf{e}_d)\,(\mathbf{e}_{d-1} + \mathbf{e}_d)\,(\mathbf{e}_{d-1})\,(2\mathbf{e}_{d-1}); \\ Q^+ &= (-2\mathbf{e}_{d-1})\,(-\mathbf{e}_{d-1})\,(-\mathbf{e}_{d-1} + \mathbf{e}_d)\,(-\mathbf{e}_{d-1} + 2\mathbf{e}_d)\,(2\mathbf{e}_d)\,(3\mathbf{e}_d); \\ Q^- &= (-2\mathbf{e}_{d-1})\,(-\mathbf{e}_{d-1})\,(-\mathbf{e}_{d-1} - \mathbf{e}_d)\,(-\mathbf{e}_{d-1} - 2\mathbf{e}_d)\,(-2\mathbf{e}_d)\,(-3\mathbf{e}_d). \end{aligned}$$

1. If $\mathbf{s}_1 \notin \{\pm\mathbf{e}_d\}$, then set $P = (-3\mathbf{s}_1)(-2\mathbf{s}_1)$ and $\mathbf{s}'_1 = D(\mathbf{s}_1)$;
   Set $P = P^+$ and $\mathbf{s}'_1 = -\mathbf{e}_{d-1}$ if $\mathbf{s}_1 = \mathbf{e}_d$ and $P = P^-$ and $\mathbf{s}'_1 = -\mathbf{e}_{d-1}$ if $\mathbf{s}_1 = -\mathbf{e}_d$;
2. If $\mathbf{s}_2 \notin \{\pm\mathbf{e}_d\}$, then set $Q = (2\mathbf{s}_2)(3\mathbf{s}_2)$ and $\mathbf{s}'_2 = D(\mathbf{s}_2)$;
   Set $Q = Q^+$ and $\mathbf{s}'_2 = -\mathbf{e}_{d-1}$ if $\mathbf{s}_2 = \mathbf{e}_d$ and $Q = Q^-$ and $\mathbf{s}'_2 = -\mathbf{e}_{d-1}$ if $\mathbf{s}_2 = -\mathbf{e}_d$;
3. Set $P[d, \mathbf{s}_1, \mathbf{s}_2] = P \diamond U(P'[d-1, \mathbf{s}'_1, \mathbf{s}'_2]) \diamond Q$:

One can check that $P[d, \mathbf{s}_1, \mathbf{s}_2]$ satisfies all three conditions of the inductive statement. ☐

Theorem 3.9 (Canonicalization)**.** *For all $d \geq 2$, $\mathsf{RQ}^d_{\mathsf{GP}}(n) \leq \mathsf{RQ}^d_{\mathsf{CGP}}(6n+1)$.*

Proof. Let $\Gamma(\mathbf{u}) = 6\mathbf{u} - \mathbf{2}$ be a map from $\mathbb{Z}^d$ to $\mathbb{Z}^d$. Given any grid-PPAD graph G^* over $\mathbb{Z}^d_n$, we now use the canonical local paths provided in Lemmas 3.7 and 3.8 to build a canonical grid-PPAD graph $G = (\mathbb{Z}^d_{6n+1}, E)$. In the procedure below, initially $E = \emptyset$:

1. [**canonicalizing the starting vertex**]: Let $\mathbf{u}^*$ be the starting vertex of G^*. Suppose $\mathbb{B}_{G^*}(\mathbf{u}^*) = (\text{“no”}, \mathbf{s})$. As $u^*_d = 1$, we have $\mathbf{s} \neq -\mathbf{e}_d$. For every edge $(\mathbf{v}_1, \mathbf{v}_2)$ appears in path $P[d, \mathbf{e}_d, \mathbf{s}]$, add $(\Gamma(\mathbf{u}^*) + \mathbf{v}_1, \Gamma(\mathbf{u}^*) + \mathbf{v}_2)$ to E;
2. [**canonicalizing the ending vertex**]: Let $\mathbf{w}^*$ be the ending vertex of G^*. Suppose $\mathbb{B}_{G^*}(\mathbf{w}^*) = (\mathbf{s}, \text{“no”})$. For every $(\mathbf{v}_1, \mathbf{v}_2)$ in $P[d, \mathbf{s}]$, add $(\Gamma(\mathbf{w}^*) + \mathbf{v}_1, \Gamma(\mathbf{w}^*) + \mathbf{v}_2)$ to E;
3. [**canonicalizing other vertices**]: For all $\mathbf{u} \in \mathbb{Z}^d_n - \{\mathbf{u}^*, \mathbf{w}^*\}$, if $\mathbb{B}_{G^*}(\mathbf{u}) = (\mathbf{s}_1, \mathbf{s}_2)$ and $\mathbf{s}_1, \mathbf{s}_2 \neq \text{“no”}$, then add $(\Gamma(\mathbf{u}) + \mathbf{v}_1, \Gamma(\mathbf{u}) + \mathbf{v}_2)$ to E for every edge $(\mathbf{v}_1, \mathbf{v}_2)$ in $P[d, \mathbf{s}_1, \mathbf{s}_2]$.

By Lemmas 3.7, 3.8 and the procedure above, $G = (\mathbb{Z}^d_{6n+1}, E)$ is a canonical grid-PPAD graph that satisfies the following two properties, from which Theorem 3.9 follows.

(C.1) The starting vertex of G is $\Gamma(\mathbf{u}^*) - 3\mathbf{e}_d$, and the ending vertex $\mathbf{w}$ of G satisfies $\|\mathbf{w} - \Gamma(\mathbf{w}^*)\| \leq 2$.

(C.2) For every vertex $\mathbf{v} \in G$, to determine $\mathbb{B}_G(\mathbf{v})$, one only need to know $\mathbb{B}_{G^*}(\mathbf{u})$ where $\mathbf{u}$ is the lexicographically smallest vertex in $\mathbb{Z}^d_n$ such that $\|\mathbf{v} - \Gamma(\mathbf{u})\| \leq 3$. ☐

3.3 From CGP^d to ZP^d: Complete the Proof of Theorem 3.2

Now we reduce CGP^d to ZP^d. The main task of this section is to, given a canonical grid-PPAD graph $G = (\mathbb{Z}^d_n, E)$ and its starting vertex $\mathbf{u}^*$, construct a discrete Brouwer function $f_G : \mathbb{Z}^d_{4n+2} \to \{\mathbf{0}\} \cup \mathbb{E}^d$ that satisfies the following two properties:

(D.1) f_G has exactly one zero point $\mathbf{r}^*$, and $\Psi^{-1}(\mathbf{r}^*)$ is the ending vertex of G, and

(D.2) For each $\mathbf{r} \in \mathbb{Z}^d_{4n+2}$, at most one query to $\mathbb{B}_G$ is needed to evaluate f_G,

where $\Psi(\mathbf{u}) = 4\mathbf{u}$ is a map from $\mathbb{Z}^d_n$ to $\mathbb{Z}^d_{4n+2}$. Immediately from these properties, we have

Theorem 3.10 (From CGP^d to ZP^d). *For all $d \geq 2$, $\mathsf{RQ}^d_{\mathsf{CGP}}(n) \leq \mathsf{RQ}^d_{\mathsf{ZP}}(4n+2)$.*

Local Geometry of Canonical Grid PPAD Graphs

To construct a Brouwer function from $G = (\mathbb{Z}^d_n, E)$, we define a set $\mathcal{I}_G \subset \mathbb{Z}^d_{4n+2}$, which looks like a collection of pipes, to embed and insulate the component of G. $\mathcal{I}_G$ has two parts, a kernel $\mathcal{K}_G$ and a boundary $\mathcal{B}_G$. We will first construct a direction-preserving function f^G on $\mathcal{B}_G$. We then extend the function onto $\mathbb{Z}^d_{4n+2}$ to define f_G. In our presentation, we will only use $\mathbf{u}, \mathbf{v}, \mathbf{w}$ to denote vertices in $\mathbb{Z}^d_n$ and use $\mathbf{p}, \mathbf{q}, \mathbf{r}$ to denote points in $\mathbb{Z}^d_{4n+2}$. Let $\mathbf{u}$ and $\mathbf{v}$ be two vertices in $\mathbb{Z}^d_n$ with $\mathbf{u} - \mathbf{v} \in \mathbb{E}^d$. We abuse $\Psi(\mathbf{uv})$ to denote the set of five integer points on line segment $\Psi(\mathbf{u})\Psi(\mathbf{v})$. Let $\mathbf{u}^*$ and $\mathbf{w}^*$ be the starting and ending vertices of G. We define

$$\mathcal{K}_G = \Big(\bigcup\nolimits_{(\mathbf{v}_1,\mathbf{v}_2)\in E} \Psi(\mathbf{v}_1\mathbf{v}_2) \Big) \cup \Big\{ \, \Psi(\mathbf{u}^*) - \mathbf{e}_d, \Psi(\mathbf{w}^*) + \mathbf{e}_d \, \Big\},$$

$$\mathcal{B}_G = \Big\{ \, \mathbf{r} \notin \mathcal{K}_G \, \Big| \, \exists \, \mathbf{r}' \in \bigcup\nolimits_{(\mathbf{v}_1,\mathbf{v}_2)\in E} \Psi(\mathbf{v}_1\mathbf{v}_2), \|\mathbf{r} - \mathbf{r}'\| = 1 \, \Big\},$$

and $\mathcal{I}_G = \mathcal{K}_G \cup \mathcal{B}_G$. For $\mathbf{u} \in \mathbb{Z}^d_n$, we use $\mathcal{C}_\mathbf{u}$ to denote $\{\, \mathbf{r} \in \mathbb{Z}^d_{4n+2}, \|\mathbf{r} - \Psi(\mathbf{u})\| \leq 2 \,\}$. As the local structure of $\mathcal{B}_G \cap \mathcal{C}_\mathbf{u}$ depends only on $\mathbb{B}_G(\mathbf{u})$, we introduce the following definitions.

Definition 3.11 (Local Kernel and Boundary). *For each pair $\pi = (\mathbf{s}_1, \mathbf{s}_2) \in \mathcal{S}^d$ with $d \geq 2$, let $\mathcal{K}_{d,\pi}$ and $\mathcal{B}_{d,\pi}$ be two subsets of $\mathbb{Z}^d_{[-2,2]} = \{-2,-1,0,1,2\}^d$, such that*

1. *if $\mathbf{s}_1 = \mathbf{s}_2 =$ "no", then $\mathcal{K}_{d,\pi} = \mathcal{B}_{d,\pi} = \emptyset$;*
2. *if $\mathbf{s}_1 =$ "no" and $\mathbf{s}_2 \neq$ "no" ($\mathbf{s}_2 = \mathbf{e}_d$), then $\mathcal{K}_{d,\pi} = \{-\mathbf{e}_d, \mathbf{0}, \mathbf{e}_d, 2\mathbf{e}_d\}$ and*
$$\mathcal{B}_{d,\pi} = \big\{ \, \mathbf{r} \in \mathbb{Z}^d_{[-2,2]} - \mathcal{K}_{d,\pi} \, \big| \, \exists \, \mathbf{r}' \in \{\mathbf{0}, \mathbf{e}_d, 2\mathbf{e}_d\}, \|\mathbf{r} - \mathbf{r}'\| = 1 \, \big\};$$
3. *if $\mathbf{s}_1 \neq$ "no" and $\mathbf{s}_2 =$ "no" ($\mathbf{s}_1 = \mathbf{e}_d$), then $\mathcal{K}_{d,\pi} = \{-2\mathbf{e}_d, -\mathbf{e}_d, \mathbf{0}, \mathbf{e}_d\}$ and*
$$\mathcal{B}_{d,\pi} = \big\{ \, \mathbf{r} \in \mathbb{Z}^d_{[-2,2]} - \mathcal{K}_{d,\pi} \, \big| \, \exists \, \mathbf{r}' \in \{-2\mathbf{e}_d, -\mathbf{e}_d, \mathbf{0}\}, \|\mathbf{r} - \mathbf{r}'\| = 1 \, \big\};$$
4. *otherwise, $\mathcal{K}_{d,\pi} = \{-2\mathbf{s}_1, -\mathbf{s}_1, 0, \mathbf{s}_2, 2\mathbf{s}_2\}$ and*
$$\mathcal{B}_{d,\pi} = \big\{ \, \mathbf{r} \in \mathbb{Z}^d_{[-2,2]} - \mathcal{K}_{d,\pi} \, \big| \, \exists \, \mathbf{r}' \in \mathcal{K}_{d,\pi}, \|\mathbf{r} - \mathbf{r}'\| = 1 \, \big\}.$$

For $\mathbf{r} \in \mathbb{Z}^d$ and set $S \subseteq \mathbb{Z}^d$, let $\mathbf{r} + S = \{\mathbf{r} + \mathbf{r}', \mathbf{r}' \in S\}$. We will use the fact that for all $\mathbf{u} \in \mathbb{Z}^d_n$, if $\pi = \mathbb{B}_G(\mathbf{u})$, then $\mathcal{K}_G \cap \mathcal{C}_\mathbf{u} = \Psi(\mathbf{u}) + \mathcal{K}_{d,\pi}$ and $\mathcal{B}_G \cap \mathcal{C}_\mathbf{u} = \Psi(\mathbf{u}) + \mathcal{B}_{d,\pi}$.

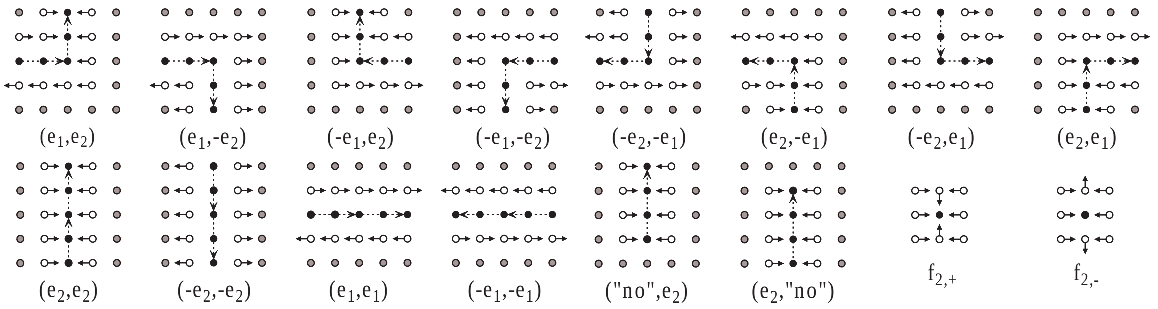

Figure 4: $f_{2,\pi}$, $f_{2,+}$ and $f_{2,-}$

The Construction of Brouwer Function

First, we define two direction-preserving functions $f_{d,+}, f_{d,-}$ from $\mathcal{B}_d$ to $\mathbb{E}^d$ for $d \geq 2$, where $\mathcal{B}_d = \{-1, 0, 1\}^d - \mathbf{0}$: For every $\mathbf{r} \in \mathcal{B}_d$, letting k be the smallest integer such that $r_k \neq 0$, $f_{d,+}(\mathbf{r}) = f_{d,-}(\mathbf{r}) = -r_k\mathbf{e}_k$ if $1 \leq k \leq d-1$ and $f_{d,+}(\mathbf{r}) = -r_d\mathbf{e}_d$, $f_{d,-}(\mathbf{r}) = r_d\mathbf{e}_d$, otherwise. Using these two functions, we inductively build a (direction-preserving) function $f_{d,\pi}$ on $\mathcal{B}_{d,\pi}$ for each $\pi = (\mathbf{s}_1, \mathbf{s}_2) \in \mathcal{S}^d$. See Figure 4 for the complete construction for $d = 2$. Informally, if $\mathbf{r} \in \mathcal{B}_{2,\pi}$ is on the left side of the "local" path, then $f_{2,\pi}(\mathbf{r}) = -\mathbf{e}_1$, otherwise it equals $\mathbf{e}_1$.

For $d \geq 3$, the construction is more complex but relatively procedural[6]. Below, we use D to denote the map $D(\mathbf{r}) = (r_1, r_2, ..., r_{d-1})$ from $\mathbb{Z}^d$ to $\mathbb{Z}^{d-1}$ and U_k to denote the map $U_k(\mathbf{r}) = (r_1, r_2, ..., r_{d-1}, k)$ from $\mathbb{Z}^{d-1}$ to $\mathbb{Z}^d$; we also extends it to sets, that is, $U_k(S) = \{U_k(\mathbf{r}), \mathbf{r} \in S\}$ for $S \subset \mathbb{Z}^{d-1}$. Let $S[k] = \{\mathbf{r} \in S, r_d = k\}$ for $S \subset \mathbb{Z}^d$.

1. **Moving within (d-1)-dimensional space:**

   When $\pi = (\mathbf{s}_1, \mathbf{s}_2)$ satisfies $\pi' = (D(\mathbf{s}_1), D(\mathbf{s}_2)) \in \mathcal{S}^{d-1}$, $\mathcal{B}_{d,\pi}$ can be decomposed into

   $$\big(U_{-1}(\mathcal{B}_{d-1,\pi'}) \cup U_{-1}(\mathcal{K}_{d-1,\pi'})\big) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup \big(U_1(\mathcal{B}_{d-1,\pi'}) \cup U_1(\mathcal{K}_{d-1,\pi'})\big)\,.$$

   We set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,\pi'}(D(\mathbf{r})))$ for $\mathbf{r} \in U_{-1}(\mathcal{B}_{d-1,\pi'}) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup U_1(\mathcal{B}_{d-1,\pi'})$.
   We set $f_{d,\pi}(\mathbf{r}) = -\mathbf{e}_{d-1}$ for $\mathbf{r} \in U_{-1}(\mathcal{K}_{d-1,\pi'})$ and $f_{d,\pi}(\mathbf{r}) = \mathbf{e}_{d-1}$ for $\mathbf{r} \in U_1(\mathcal{K}_{d-1,\pi'})$.

2. **Moving along $\pm\mathbf{e}_d$:** In this case, we will use the fact $D(\mathbf{r}) \in \mathcal{B}_{d-1}$, for all $\mathbf{r} \in \mathcal{B}_{d,\pi}$.

   When $\pi = (\mathbf{e}_d, \mathbf{e}_d)$, ("no", $\mathbf{e}_d$) or ($\mathbf{e}_d$, "no"), $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,+}(D(\mathbf{r})))$ for all $\mathbf{r} \in \mathcal{B}_{d,\pi}$.
   When $\pi = (-\mathbf{e}_d, -\mathbf{e}_d)$, we set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,-}(D(\mathbf{r})))$ for all $\mathbf{r} \in \mathcal{B}_{d,\pi}$.

3. **Moving between $\mathbf{e}_{d-1}$ and $\pm\mathbf{e}_d$:**

   (a) When $\pi = (\mathbf{s}_1, \mathbf{s}_2) = (\mathbf{e}_{d-1}, \mathbf{e}_d)$, let $\pi' = (D(\mathbf{s}_1), \text{"no"}) \in \mathcal{S}^{d-1}$. We have

   $$\mathcal{B}_{d,\pi}[2] = U_2(\mathcal{B}_{d-1}),\ \mathcal{B}_{d,\pi}[1] = U_1(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}) - \{\mathbf{e}_d\}, \mathcal{B}_{d,\pi}[-2] = \emptyset,$$
   $$\mathcal{B}_{d,\pi}[0] = U_0(\mathcal{B}_{d-1,\pi'}) \cup \{\mathbf{e}_{d-1}\},\ \mathcal{B}_{d,\pi}[-1] = U_{-1}(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}).$$

   We set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,+}(D(\mathbf{r})))$ for $\mathbf{r} \in \mathcal{B}_{d,\pi}[2]$;

---

[6] Sorry for so many cases. You will find that they are progressively easier to understand.

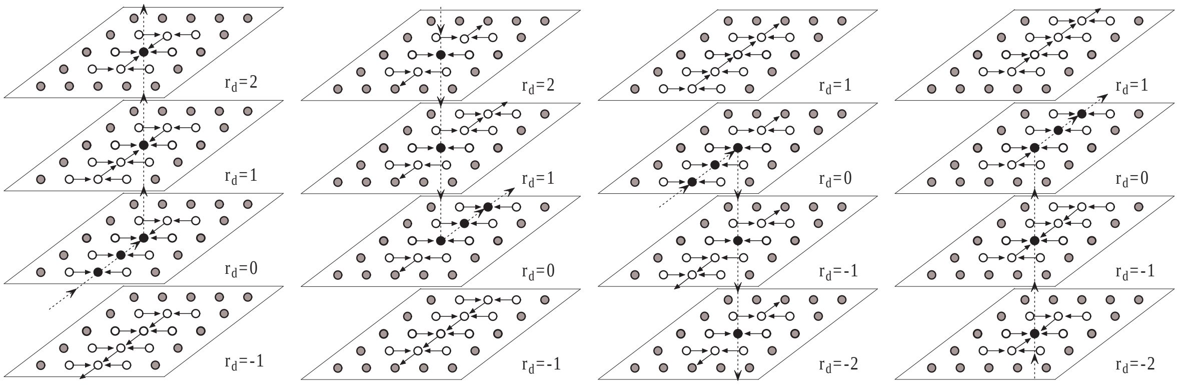

Figure 5: $f_{3,\pi}$ where $\pi = (\mathbf{e}_{d-1}, \mathbf{e}_d), (-\mathbf{e}_d, \mathbf{e}_{d-1}), (\mathbf{e}_{d-1}, -\mathbf{e}_d)$ and $(\mathbf{e}_d, \mathbf{e}_{d-1})$

$f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,\pi'}(D(\mathbf{r})))$ for $\mathbf{r} \in U_1(\mathcal{B}_{d-1,\pi'}) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup U_{-1}(\mathcal{B}_{d-1,\pi'})$;
$f_{d,\pi}(\mathbf{e}_{d-1} + \mathbf{e}_d) = f_{d,\pi}(\mathbf{e}_{d-1}) = -\mathbf{e}_{d-1}$; $f_{d,\pi}(\mathbf{r}) = -\mathbf{e}_{d-1}$ for $\mathbf{r} \in U_{-1}(\mathcal{K}_{d-1,\pi'})$;
$f_{d,\pi}(\mathbf{r}) = \mathbf{e}_{d-1}$ for $\mathbf{r} \in U_1(\mathcal{K}_{d-1,\pi'}) - \{\mathbf{e}_d, \mathbf{e}_{d-1} + \mathbf{e}_d\}$.

(b) When $\pi = (\mathbf{s}_1, \mathbf{s}_2) = (-\mathbf{e}_d, \mathbf{e}_{d-1})$, let $\pi' = (\text{“no”}, D(\mathbf{s}_2)) \in \mathcal{S}^{d-1}$. We have

$$\mathcal{B}_{d,\pi}[-2] = \emptyset,\ \mathcal{B}_{d,\pi}[2] = U_2(\mathcal{B}_{d-1}),\ \mathcal{B}_{d,\pi}[1] = U_1(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}) - \{\mathbf{e}_d\},$$
$$\mathcal{B}_{d,\pi}[0] = U_0(\mathcal{B}_{d-1,\pi'}) \cup \{-\mathbf{e}_{d-1}\},\ \mathcal{B}_{d,\pi}[-1] = U_{-1}(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}).$$

We set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,-}(D(\mathbf{r})))$ for $\mathbf{r} \in \mathcal{B}_{d,\pi}[2]$;
$f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,\pi'}(D(\mathbf{r})))$ for $\mathbf{r} \in U_1(\mathcal{B}_{d-1,\pi'}) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup U_{-1}(\mathcal{B}_{d-1,\pi'})$;
$f_{d,\pi}(\mathbf{r}) = -\mathbf{e}_{d-1}$ for $\mathbf{r} \in U_{-1}(\mathcal{K}_{d-1,\pi'})$; $f_{d,\pi}(-\mathbf{e}_{d-1} + \mathbf{e}_d) = f_{d,\pi}(-\mathbf{e}_{d-1}) = -\mathbf{e}_{d-1}$;
$f_{d,\pi}(\mathbf{r}) = \mathbf{e}_{d-1}$ for $\mathbf{r} \in U_1(\mathcal{K}_{d-1,\pi'}) - \{-\mathbf{e}_{d-1} + \mathbf{e}_d, \mathbf{e}_d\}$.

(c) When $\pi = (\mathbf{s}_1, \mathbf{s}_2) = (\mathbf{e}_{d-1}, -\mathbf{e}_d)$, let $\pi' = (D(\mathbf{s}_1), \text{“no”}) \in \mathcal{S}^{d-1}$. We have

$$\mathcal{B}_{d,\pi}[2] = \emptyset,\ \mathcal{B}_{d,\pi}[1] = U_1(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}),\ \mathcal{B}_{d,\pi}[0] = U_0(\mathcal{B}_{d-1,\pi'}) \cup \{\mathbf{e}_{d-1}\},$$
$$\mathcal{B}_{d,\pi}[-1] = U_{-1}(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}) - \{-\mathbf{e}_d\},\ \mathcal{B}_{d,\pi}[-2] = U_{-2}(\mathcal{B}_{d-1}).$$

We set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,-}(D(\mathbf{r})))$ for $\mathbf{r} \in \mathcal{B}_{d,\pi}[-2]$;
$f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,\pi'}(D(\mathbf{r})))$ for $\mathbf{r} \in U_1(\mathcal{B}_{d-1,\pi'}) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup U_{-1}(\mathcal{B}_{d-1,\pi'})$;
$f_{d,\pi}(\mathbf{r}) = \mathbf{e}_{d-1}$ for $\mathbf{r} \in U_1(\mathcal{K}_{d-1,\pi'})$; $f_{d,\pi}(\mathbf{e}_{d-1}) = f_{d,\pi}(\mathbf{e}_{d-1} - \mathbf{e}_d) = \mathbf{e}_{d-1}$;
$f_{d,\pi}(\mathbf{r}) = -\mathbf{e}_{d-1}$ for $\mathbf{r} \in U_{-1}(\mathcal{K}_{d-1,\pi'}) - \{-\mathbf{e}_d, \mathbf{e}_{d-1} - \mathbf{e}_d\}$.

(d) When $\pi = (\mathbf{s}_1, \mathbf{s}_2) = (\mathbf{e}_d, \mathbf{e}_{d-1})$, let $\pi' = (\text{“no”}, D(\mathbf{s}_2)) \in \mathcal{S}^{d-1}$. We have

$$\mathcal{B}_{d,\pi}[2] = \emptyset,\ \mathcal{B}_{d,\pi}[1] = U_1(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}),\ \mathcal{B}_{d,\pi}[0] = U_0(\mathcal{B}_{d-1,\pi'}) \cup \{-\mathbf{e}_{d-1}\},$$
$$\mathcal{B}_{d,\pi}[-1] = U_{-1}(\mathcal{B}_{d-1,\pi'} \cup \mathcal{K}_{d-1,\pi'}) - \{-\mathbf{e}_d\},\ \mathcal{B}_{d,\pi}[-2] = U_{-2}(\mathcal{B}_{d-1}).$$

We set $f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,+}(D(\mathbf{r})))$ for $\mathbf{r} \in \mathcal{B}_{d,\pi}[-2]$;
$f_{d,\pi}(\mathbf{r}) = U_0(f_{d-1,\pi'}(D(\mathbf{r})))$ for $\mathbf{r} \in U_1(\mathcal{B}_{d-1,\pi'}) \cup U_0(\mathcal{B}_{d-1,\pi'}) \cup U_{-1}(\mathcal{B}_{d-1,\pi'})$;
$f_{d,\pi}(\mathbf{r}) = \mathbf{e}_{d-1}$ for $\mathbf{r} \in U_1(\mathcal{K}_{d-1,\pi'})$; $f_{d,\pi}(-\mathbf{e}_{d-1}) = f_{d,\pi}(-\mathbf{e}_{d-1} - \mathbf{e}_d) = \mathbf{e}_{d-1}$;
$f_{d,\pi}(\mathbf{r}) = -\mathbf{e}_{d-1}$ for $\mathbf{r} \in U_{-1}(\mathcal{K}_{d-1,\pi'})) - \{-\mathbf{e}_d, -\mathbf{e}_{d-1} - \mathbf{e}_d\}$.

$f_G(\mathbf{r})$, where $\mathbf{r} \in \mathbb{Z}^d_{4n+2}$

```
let $\mathbf{u}^*$ and $\mathbf{w}^*$ be the starting and ending vertices of $G$, $\mathbf{p}^* = \Psi(\mathbf{u}^*)$ and $\mathbf{q}^* = \Psi(\mathbf{w}^*)$
if $r_d = 1$ or 2, and $D(\mathbf{r}) = D(\mathbf{p}^*)$ then $f_G(\mathbf{r}) = \mathbf{e}_d$
else if $r_d = 1$ or 2, and $\|D(\mathbf{r}) - D(\mathbf{p}^*)\| = 1$ then
    let $k$ denote the smallest integer such that $r_k \neq r^*_k$, $f_G(\mathbf{r}) = (p^*_k - r_k)\mathbf{e}_k$
else if $\mathbf{r} = \mathbf{q}^*$ then $f_G(\mathbf{r}) = 0$
else if $\mathbf{r} = \mathbf{q}^* + \mathbf{e}_d$ then $f_G(\mathbf{r}) = -\mathbf{e}_d$
else if $\mathbf{r} \in \mathcal{K}_G$ then $f_G(\mathbf{r}) = \mathbf{e}_d$
else if $\mathbf{r} \in \mathcal{B}_G$ then $f_G(\mathbf{r}) = f^G(\mathbf{r})$
else if $r_d = 1$ (and $\|D(\mathbf{r}) - D(\mathbf{p}^*)\| \geq 2$) then
    let $k$ denote the smallest integer such that $r_k \neq r^*_k$, $f_G(\mathbf{r}) = \text{sign}(p^*_k - r_k)\mathbf{e}_k$
else $f_G(\mathbf{r}) = -\mathbf{e}_d$
```

Figure 6: Construction of Function f_G from f^G

Lemma 3.12 (Locally Directional Preserving)**.** *For every $\pi \in \mathcal{S}^d$, $f_{d,\pi}$ is direction-preserving on $\mathcal{B}_{d,\pi}$.*

Proof. We prove the lemma by induction on d. The base case when $d = 2$ is trivial. We now consider the case when $d > 2$ and assume inductively that the statement is true for $d - 1$.

First, $\pi = (\mathbf{e}_d, \mathbf{e}_d)$, ("no", $\mathbf{e}_d$), ($\mathbf{e}_d$, "no") or $(-\mathbf{e}_d, -\mathbf{e}_d)$. The statement follows from the fact that $f_{d-1,+}$ and $f_{d-1,-}$ are direction-preserving on $\mathcal{B}_{d-1}$. Second, $\pi = (\mathbf{s}_1, \mathbf{s}_2)$ satisfies $\pi' = (D(\mathbf{s}_1), D(\mathbf{s}_2)) \in \mathcal{S}^{d-1}$. By the inductive hypothesis, $f_{d-1,\pi'}$ is direction-preserving, from which the statement follows. Third, $\pi = (\mathbf{e}_{d-1}, \mathbf{e}_d)$, $(\mathbf{e}_{d-1}, -\mathbf{e}_d)$, $(-\mathbf{e}_d, \mathbf{e}_{d-1})$ or $(\mathbf{e}_d, \mathbf{e}_{d-1})$. One can prove the following statement by induction on d.

> For $\pi_1 = $ ("no", $\mathbf{e}_d$) and $\pi_2 = $ ($\mathbf{e}_d$, "no"), $\mathcal{B}_{d,\pi_1} \cap \mathcal{B}_d = \mathcal{B}_{d,\pi_2} \cap \mathcal{B}_d = \mathcal{B}_d - \{-\mathbf{e}_d, \mathbf{e}_d\}$. Moreover for each $\mathbf{r} \in \mathcal{B}_d - \{-\mathbf{e}_d, \mathbf{e}_d\}$, $f_{d,\pi_1}(\mathbf{r}) = f_{d,\pi_2}(\mathbf{r}) = f_{d,+}(\mathbf{r}) = f_{d,-}(\mathbf{r})$.

To show $f_{d,\pi}$ is direction-preserving on $\mathcal{B}_{d,\pi}$, it suffices to check $\|\mathbf{r}_1 - \mathbf{r}_2\| > 1$, for all pairs $\mathbf{r}_1, \mathbf{r}_2 \in \mathcal{B}_{d,\pi}$ such that$f_{d,\pi}(\mathbf{r}_1) = \mathbf{e}_{d-1}$ and $f_{d,\pi}(\mathbf{r}_2) = -\mathbf{e}_{d-1}$. □

With these local functions $f_{d,\pi}$, we can build a global function f^G from $\mathcal{B}_G$ to $\{\pm\mathbf{e}_1, ..., \pm\mathbf{e}_{d-1}\}$ as following: for every $\mathbf{r} \in \mathcal{B}_G$, we set $f^G(\mathbf{r}) = f_{d,\pi}(\mathbf{r} - \Psi(\mathbf{u}))$, where $\mathbf{u}$ is the lexicographically smallest vertex in $\mathbb{Z}^d_n$ such that $\mathbf{r} \in \mathcal{C}_{\mathbf{u}}$ and $\pi = \mathbb{B}_G(\mathbf{u})$,

Lemma 3.13. *For every canonical grid PPAD graph G over $\mathbb{Z}^d_n$, f^G is direction-preserving on set $\mathcal{B}_G$.*

Proof. By Lemma 3.12, it suffices to prove the following: For $\mathbf{r} \in \mathcal{B}_G$, if $\mathbf{r} \in \mathcal{C}_{\mathbf{u}} \cap \mathcal{C}_{\mathbf{v}}$ where $\mathbf{u}$, $\mathbf{v} \in \mathbb{Z}^d_n$, then $f_{d,\pi_1}(\mathbf{r} - \Psi(\mathbf{u})) = f_{d,\pi_2}(\mathbf{r} - \Psi(\mathbf{v}))$, where $\pi_1 = \mathbb{B}_G(\mathbf{u})$ and $\pi_2 = \mathbb{B}_G(\mathbf{v})$. We will use the fact that $\mathbf{s} = \mathbf{u} - \mathbf{v} \in \mathbb{E}^d$ and either $(\mathbf{u}, \mathbf{v}) \in G$ or $(\mathbf{v}, \mathbf{u}) \in G$.

For $S \subset \mathbb{Z}^d$ and $\mathbf{p} \in \mathbb{Z}^d$, we use $S + \mathbf{p}$ to denote $\{\mathbf{r} \in \mathbb{Z}^d \mid \mathbf{r} = \mathbf{r}' + \mathbf{p}, \mathbf{r}' \in S\}$. The lemma is a direct consequence of the following statement which can be proved by induction on d.

> For all $\mathbf{s} = b\mathbf{e}_k \in \mathbb{E}^d$ with $b \in \{\pm 1\}$ and $k \in [1 : d]$, if $\pi_1 = (\mathbf{s}_1, \mathbf{s}) \in \mathcal{S}^d$ and $\pi_2 = (\mathbf{s}, \mathbf{s}_2) \in \mathcal{S}^d$ then $\{\mathbf{r} \in \mathcal{B}_{d,\pi_1}, r_k = 2b\} = \{\mathbf{r} \in \mathcal{B}_{d,\pi_2}, r_k = -2b\} + 4\mathbf{s}$, and for every $\mathbf{r}$ in the former set, $f_{d,\pi_1}(\mathbf{r}) = f_{d,\pi_2}(\mathbf{r} - 4\mathbf{s})$. □

Finally, to extend f^G onto $\mathbb{Z}^d_{4n+2}$ to define our function f_G, we apply the procedure given in Fig. 6. It is somewhat tedious but procedural to check that f_G satisfies both **Property D.2** and **D.1** stated at the beginning of this subsection.

4 Randomized Lower Bound for ES^d

The technical objective of this section is to construct a distribution $\mathcal{S}$ of d-non-repeating strings and show that, for a random string S drawn according to $\mathcal{S}$, every deterministic algorithm for ES^d needs expected $(\Omega(n))^d$ queries to $\mathbb{B}_S$. Thus, by Yao's Minimax Principle [32], we have $\mathsf{RQ}^d_{\mathsf{ES}}(n) = (\Omega(n))^d$. Our main Theorem 2.2 then follows from Theorems 3.1 and 3.2.

We apply random permutations hierarchically to define distribution $\mathcal{S}$ to ensure that a random string from $\mathcal{S}$ has sufficient entropy that its search problem is expected to be difficult. The use the hierarchical structure guarantees that each string in $\mathcal{S}$ is d-non-repeating.

4.1 Hierarchical Construction of Random d-Non-Repeating Strings

We first define our hierarchical framework. Let $\mathbb{J}_n = [2 : 2n+2]$, $\mathbb{O}_n = \{3, 5, ..., 2n+1\}$ and $\mathbb{F}_n = \{4, 6, ..., 2n+2\}$. Let $S_0 = 2$, $S_1 = 3 \circ 4$, ..., $S_n = (2n+1) \circ (2n+2)$. Each permutation π from $[1 : n]$ to $[1 : n]$ defines a string $C = S_0 \circ S_{\pi(1)} \circ \cdots \circ S_{\pi(n)}$ which we refer to as a *connector* over $\mathbb{J}_n$.

Let $r[C] = 2\pi(n) + 2$, the last symbol of C. We use $\phi_C(2)$ to denote the right neighbor of 2. Each $s \in \mathbb{J}_n - \{2, r[C_\pi]\}$ has two neighbors in C. The left neighbor of an even s is $s - 1$, we use $\phi_C(s)$ to denote its right neighbor; the right neighbor of an odd s is $s + 1$, and we use $\phi_C(s)$ to denote its left neighbor. Clearly, if $\phi_C(s) = t$ then $\phi_C(t) = s$.

Our hierarchical framework is built on $T_{n,d}$, the rooted complete-$(2n+1)$-nary tree of height d. In $T_{n,d}$, each internal node u is connected to its $(2n+1)$ children by edges with distinct labels from $\mathbb{J}_n$; if u is connected to v by an edge labeled with j, then we call v the j^{th}*-successor* of u. Each node v of $T_{n,d}$ has a *natural name*, $\mathsf{name}(v)$, the concatenation of labels along the path from the root of $T_{n,d}$ to v. Let $\mathsf{height}(v)$ and $\mathsf{level}(v)$ denote the height and level of node v in the tree. For example, the height of the root is d and the level of the root is 0.

Definition 4.1 (Tree-of-Connectors). *An (n, d)-ToC $\mathcal{T}$ is a tree $T_{n,d}$ in which each internal node v is associated with a connector C_v over $\mathbb{J}_n$. The $r[C_v]^{th}$-successor is referred to as the* last successor *of v. The* tail *of v, $\mathsf{tail}(v)$, is the leaf reachable from v by last-successor relations. The tail of a leaf is itself. The* tail *of $\mathcal{T}$, $\mathsf{tail}(\mathcal{T})$, is the tail of its root. The* head *of a leaf u, $\mathsf{head}(u)$, is the ancestor of u with the largest height such that u is its tail.*

Definition 4.2 (Valid ToC)**.** *An* (n,d)-ToC $\mathcal{T}$ *is* valid *if for all internal* v *and for each pair of* $s,t \in \mathbb{J}_n$ *with* $\phi_{C_v}(s) = t$, $\mathsf{name}\,(u_s)$ *and* $\mathsf{name}\,(u_t)$ *share a common suffix of length* $\mathsf{height}\,(v) - 1$, *where* u_s *and* u_t*, respectively, are the tails of the* s^{th}*-successor and* t^{th}*-successor of* v*.*

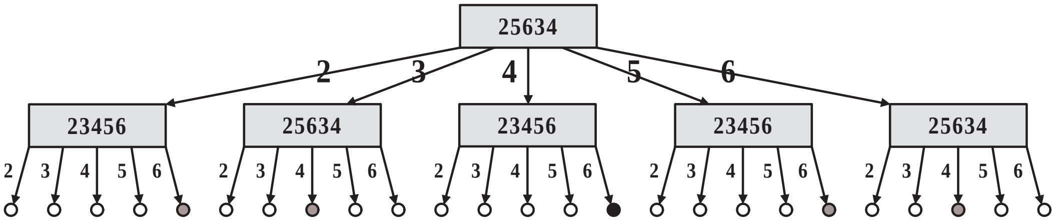

Figure 7: A valid $(2,2)$-ToC $\mathcal{T}$

Definition 4.3 ($\mathbb{B}_{\mathcal{T}}$ for accessing $\mathcal{T}$)**.** *Suppose* $\mathcal{T}$ *is a valid* (n,d)*-ToC. The input to* $\mathbb{B}_{\mathcal{T}}$ *is a point* $\mathbf{q}$ *from* $(\mathbb{J}_n)^d$ *(defining the name of a leaf* u *in* $\mathcal{T}$*). Let* $h = \mathsf{height}\,(\mathsf{head}\,(u))$*. If* u *is the tail of* $\mathcal{T}$*, i.e.,* $h = d$*, then* $\mathbb{B}_{\mathcal{T}} = \mathcal{T}$*. Otherwise, let* $v_1 = \mathsf{head}\,(u)$ *and let* v *be the parent of* v_1*. Note that* v_1 *is the* q_{d-h}^{th}*-successor of* v*. Let* $\mathcal{T}_1$ *be the tree rooted at* v_1*. As* $u \neq \mathsf{tail}\,(v)$*,* $\phi_{C_v}(q_{d-h})$ *is defined and let* $\mathcal{T}_2$ *be the subtree rooted the* $\phi_{C_v}(q_{d-h})^{th}$*-successor of* v*. Then,* $\mathbb{B}_{\mathcal{T}}(\mathbf{q}) = (h, \phi_{C_v}(q_{d-h}), \mathcal{T}_1, \mathcal{T}_2)$*.*

We now define our final search problem $\mathsf{Name\text{-}the\text{-}Tail}$, on a valid (n,d)-ToC. The search problem NT^d is: Given a valid (n,d)-ToC $\mathcal{T}^*$ accessible by $\mathbb{B}_{\mathcal{T}^*}$, find the name of its tail. We will prove Theorem 4.4 in Section 4.3. Below, we prove Theorem 4.5 to reduce NT^d to ES^d.

Theorem 4.4 (Complexity of NT^d)**.** *For all sufficiently large* n*,*

$$\mathsf{RQ}^d_{\mathsf{NT}}(n) \geq \frac{1}{2}\left(\frac{n}{2 \cdot 24^d}\right)^d.$$

Theorem 4.5 (From NT^d to ES^d)**.** *For all* $d \geq 1$*,* $\mathsf{RQ}^d_{\mathsf{NT}}(n) \leq \mathsf{RQ}^d_{\mathsf{ES}}(4n+4)$*.*

Proof. We need to build a d-non-repeating string from a valid (n,d)-ToC $\mathcal{T}$. In fact, we will construct two strings $S[\mathcal{T}]$ and $Q[\mathcal{T}]$ over $\mathbb{Z}_{4n+4}$, each has length $\Theta(n^d) \cdot d$. $S[\mathcal{T}]$ starts with $\mathbf{s}_d$ and ends with $\mathcal{F}(\mathsf{name}\,(\mathsf{tail}\,(\mathcal{T}))) \in \mathbb{Z}^d$ while $Q[\mathcal{T}]$ starts with $\mathcal{F}(\mathsf{name}\,(\mathsf{tail}\,(\mathcal{T})))$ and ends with $\mathbf{s}_d$, where for $\mathbf{p} \in \mathbb{Z}^d$ $\mathcal{F}(\mathbf{p}) = (2p_1, ..., 2p_{d-1}, 2p_d - 1)$ and $\mathbf{s}_d \in \mathbb{Z}^d$ defined to be $s_1 = 1$, and $\mathbf{s}_d = (2, ..., 2, 1)$ for $d > 1$.

For any two strings $S_1 = a_1...a_k$ and $S_2 = b_1...b_t$, let $S_1 \circ S_2 = a_1...a_k b_1...b_t$. For $d \geq 1$, if $a_{k-d+i} = b_i$ for all $1 \leq i \leq d$, then let $S_1 \circ_d S_2 = a_1...a_k b_{d+1}...b_t$. Given a string S over $\mathbb{Z}$ of length $k \cdot d$, we write S as $\mathbf{u}_1...\mathbf{u}_k$ with $\mathbf{u}_i \in \mathbb{Z}^d$. Let $\mathsf{insert}_d\,(S,t) = \mathbf{u}_1 \circ t \circ \mathbf{u}_2 \circ t ... \mathbf{u}_{k-1} \circ t \circ \mathbf{u}_k$.

We use the following recursive procedure. Let r be the root of $\mathcal{T}$. Assume $C_r = a_1...a_{2n+1}$. When $d = 1$, we set $S[\mathcal{T}] = 1b_1b_2...b_{2n+1}$ and $Q[\mathcal{T}] = b_{2n+1}...b_2b_11$, where $b_i = 2a_i - 1$. When $d \geq 2$,

1. let $\mathcal{T}_i$ be the subtree of $\mathcal{T}$ rooted at the a_i^{th}-successor of r and let $\mathbf{p}_i \in (\mathbb{F}_n)^{d-1}$ be the name of the tail of $\mathcal{T}_i$ given by $\mathcal{T}_i$ (not by $\mathcal{T}$).

2. for every odd $i \in [1:2n{+}1]$, set $S'_i = \mathsf{insert}_{d-1}(S[\mathcal{T}_i], 2a_i)$ (which starts with $\mathbf{s}_{d-1}$ and ends with $\mathcal{F}(\mathbf{p}_i)$) and for every even $i \in [1:2n+1]$ set $S'_i = \mathsf{insert}_{d-1}(Q[\mathcal{T}_i], 2a_i)$ (which starts with $\mathcal{F}(\mathbf{p}_i)$ and ends with $\mathbf{s}_{d-1}$); $S[\mathcal{T}] = \mathbf{s}_d \circ_{d-1} S'_1 \circ_{d-1} S'_2 \circ_{d-1} S'_3 \circ_{d-1} ... \circ_{d-1} S'_{2n} \circ_{d-1} S'_{2n+1}$;
3. for every odd $i \in [1:2n+1]$, set $Q'_i = \mathsf{insert}_{d-1}(Q[\mathcal{T}_i], 2a_i)$ (which starts with $\mathcal{F}(\mathbf{p}_i)$ and ends with $\mathbf{s}_{d-1}$) and for every even $i \in [1:2n+1]$, set $Q'_i = \mathsf{insert}_{d-1}(S[\mathcal{T}_i], 2a_i)$ (which starts with $\mathbf{s}_{d-1}$ and ends with $\mathcal{F}(\mathbf{p}_i)$); $Q[\mathcal{T}] = (2a_{2n+1}) \circ Q'_{2n+1} \circ_{d-1} Q'_{2n} \circ_{d-1} ... \circ_{d-1} Q'_2 \circ_{d-1} Q'_1 \circ \mathbf{s}_d$.

The two strings for the example in Figure 7 above are:

$$S[\mathcal{T}] = {}^{2}{}_{1}{}^{4}{}_{3}{}^{4}{}_{5}{}^{4}{}_{7}{}^{4}{}_{9}{}^{4}{}_{11}{}^{10}{}_{9}{}^{10}{}_{7}{}^{10}{}_{5}{}^{10}{}_{3}{}^{10}{}_{1}{}^{12}{}_{3}{}^{12}{}_{9}{}^{12}{}_{11}{}^{12}{}_{5}{}^{12}{}_{7}{}^{6}{}_{5}{}^{6}{}_{11}{}^{6}{}_{9}{}^{6}{}_{3}{}^{6}{}_{1}{}^{8}{}_{3}{}^{8}{}_{5}{}^{8}{}_{7}{}^{8}{}_{9}{}^{8}{}_{11}$$

$$Q[\mathcal{T}] = {}^{8}{}_{11}{}^{8}{}_{9}{}^{8}{}_{7}{}^{8}{}_{5}{}^{8}{}_{3}{}^{8}{}_{1}{}^{6}{}_{3}{}^{6}{}_{9}{}^{6}{}_{11}{}^{6}{}_{5}{}^{6}{}_{7}{}^{12}{}_{5}{}^{12}{}_{11}{}^{12}{}_{9}{}^{12}{}_{3}{}^{12}{}_{1}{}^{10}{}_{3}{}^{10}{}_{5}{}^{10}{}_{7}{}^{10}{}_{9}{}^{10}{}_{11}{}^{4}{}_{9}{}^{4}{}_{7}{}^{4}{}_{5}{}^{4}{}_{3}{}^{4}{}_{1}{}^{2}{}_{1}$$

The correctness of our construction can be established using the next two lemmas. ☐

Lemma 4.6 (Non-Repeating)**.** *If $\mathcal{T}$ is a valid (n,d)-*ToC*, then both $S[\mathcal{T}]$ and $Q[\mathcal{T}]$ are d-non-repeating.*

Proof. We prove the lemma by induction on d. The base case when $d = 1$ is trivial.

Assume $d > 1$ and also inductively that the statement is true for $d-1$. Suppose for the sake of contradiction that $S' = a_1 a_2 ... a_d \in \mathbb{Z}^d_{4n+4}$ appears in $S[\mathcal{T}]$ more than once. Note that exactly one symbol in S', say a_t, is odd. Let $k \in [1:d]$ be the following integer: if $t = d$, then $k = 1$; otherwise, $k = t+1$.

First, if $a_k = 2$, then S' appears in $S[\mathcal{T}]$ implies $(a_1, a_2, ..., a_d) = \mathbf{s}_d$. But such S' only appears in $S[\mathcal{T}]$ once, which contradicts with the assumption.

Otherwise, let $S'' \in \mathbb{Z}^{d-1}_{4n+4}$ be the string obtained by removing a_k from S', then: if $a_k/2$ is odd, then S'' appears in $Q[\mathcal{T}_{a_k/2}]$ more than once; otherwise, S'' appears in $S[\mathcal{T}_{a_k/2}]$ more than once, which contradicts with the inductive hypothesis.

The proof for string $Q[\mathcal{T}]$ is similar. ☐

Lemma 4.7 (Asking $\mathbb{B}_{\mathcal{T}}$)**.** *Suppose $\mathcal{T}$ is a valid (n,d)-ToC and $S = S[\mathcal{T}]$ and $Q = Q[\mathcal{T}]$. For any $\mathbf{u} \in \mathbb{Z}^d_{4n+4}$, we can compute $\mathbb{B}_S(\mathbf{u})$ and $\mathbb{B}_Q(\mathbf{u})$ by querying $\mathbb{B}_{\mathcal{T}}$ at most once.*

Proof. We need the following two propositions. Proposition 4.9 can be proved by mathematical induction on d.

Proposition 4.8 (Vectors not in S and Q)**.** *Let $V_1 = \{2,4,...,4n+4\}^{d-1} \times \{1,3,...,4n+3\}$ and $V_k = \{2,4,...,4n+4\}^{k-2} \times \{1,3,...,4n+3\} \times \{2,4,...,4n+4\}^{d-k+1}$ for $2 \le k \le d$. If $\mathbf{u} \notin \cup_k V_k$ then $\mathbf{u}$ neither appears in S nor in Q.*

Proposition 4.9 (All the same)**.** *Let $\mathcal{T}$ and $\mathcal{T}'$ two valid (n,d)-ToCs. If $\mathbf{u} \in \cup_k V_k$ and $u_i = 1$ or $u_i = 2$ for some $1 \le i \le d$, then $\mathbb{B}_{S[\mathcal{T}]}(\mathbf{u}) = \mathbb{B}_{S[\mathcal{T}']}(\mathbf{u})$ and $\mathbb{B}_{Q[\mathcal{T}]}(\mathbf{u}) = \mathbb{B}_{Q[\mathcal{T}']}(\mathbf{u})$.*

We first consider two simple cases for which we don't even need to query $\mathbb{B}_{\mathcal{T}}$.

1. When $\mathbf{u} \notin \cup_k V_k$, by Proposition 4.8, $\mathbb{B}_{S[\mathcal{T}]}(\mathbf{u}) = \mathbb{B}_{Q[\mathcal{T}]}(\mathbf{u}) = $ ("no","no");

2. When $\mathbf{u} \in \cup_k V_k$ and $u_i = 1$ or $u_i = 2$ for some $1 \le i \le d$, by Proposition 4.9, we can compute $\mathbb{B}_S(\mathbf{u})$ and $\mathbb{B}_Q(\mathbf{u})$ from the valid (n,d)-tree in which every connector is generated by the identity permutation from $\{1,\dots,n\}$ to $\{1,\dots,n\}$.

Now we can assume $\mathbf{u} \in \cup_k U_k$ where $U_1 = \{4,6,...,4n+4\}^{d-1} \times \{3,5,...,4n+3\}$ and $U_k = \{4,6...,4n+4\}^{k-2} \times \{3,5,...,4n+3\} \times \{4,6,...,4n+4\}^{d-k+1}$ for $2 \le k \le d$. First note that there is exactly one odd entry in $\mathbf{u}$. If $\mathbf{u} \in U_k$ then let $\mathbf{u}'$ be the string obtained from $\mathbf{u}$ by $k-1$ left-rotations. Not the last entry of $\mathbf{u}'$ is odd. Let $\mathbf{q}$ be the vector in $(\mathbb{J}_n)^d$ where $q_i = u'_i/2$ for $1 \le i \le d-1$ and $q_d = (u'_d+1)/2$. We now prove a stronger statement which implies that $\mathbb{B}_S(\mathbf{u})$ and $\mathbb{B}_Q(\mathbf{u})$ can be computed from $\mathbb{B}_{\mathcal{T}}(\mathbf{q})$.

> *For all $\mathbf{q} \in (\mathbb{J}_n)^d$, we can determine $\mathbb{B}_S(\mathbf{u}_i)$ and $\mathbb{B}_Q(\mathbf{u}_i)$ from $\mathbb{B}_{\mathcal{T}}(\mathbf{q})$, where $\mathbf{u}_k$ is the vector in U_k obtained from $\mathcal{F}(\mathbf{q})$ by $k-1$ right rotations, for $1 \le k \le d$.*

If $\mathbb{B}_{\mathcal{T}}(\mathbf{q}) = \mathcal{T}$, the statement is clearly true. Otherwise, assuming $\mathbb{B}_{\mathcal{T}}(\mathbf{q}) \ne \mathcal{T}$, we prove the statement by induction on d. The base case when $d=1$ is trivial. For $d \ge 2$, let $\mathbf{q}' = (q_2, q_3, ..., q_d)$ and let $\mathbf{v}_k$ be the vector generated from $\mathcal{F}(\mathbf{q}')$ by $k-1$ right rotations, for $1 \le k \le d-1$. Let $\mathcal{T}'$ be the subtree of $\mathcal{T}$ rooted the q_1-successor of the root of $\mathcal{T}$. As $\mathbb{B}_{\mathcal{T}'}(\mathbf{q}')$ is contained in $\mathbb{B}_{\mathcal{T}}(\mathbf{q})$, we can determine $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_i)$ and $\mathbb{B}_{Q[\mathcal{T}']}(\mathbf{v}_i)$ for using our inductive hypothesis, from which, we will show below, we can determine $\mathbb{B}_S(\mathbf{u}_i)$ and $\mathbb{B}_Q(\mathbf{u}_i)$.

We will only prove the case for $\mathbb{B}_S(\mathbf{u}_i)$ when q_1 is even. All other cases are similar.

Note that the first entry of $\mathbf{v}_i$ is not 2, so for all $1 \le i \le d-1$, $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_i) \ne (\text{“no”}, a)$. Also, for $i > 1$ the last entry of $\mathbf{v}_i$ is even, so, $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_i) \ne (a, \text{“no”})$. Therefore, for all $i \in [2 : d-1]$, $\mathbb{B}_S(\mathbf{u}_i) = \mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_i)$. For $i = 1$ or d, if $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_1) = (\text{“no”}, \text{“no”})$, then $\mathbb{B}_S(\mathbf{u}_i) = (\text{“no”}, \text{“no”})$; if $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_1) = (a,b)$, then $\mathbb{B}_S(\mathbf{u}_1) = (a, 2q_1)$, $\mathbb{B}_S(\mathbf{u}_d) = (2q_1, b)$; if $\mathbb{B}_{S[\mathcal{T}']}(\mathbf{v}_1) = (a, \text{“no”})$, letting the second component of $\mathbb{B}_{\mathcal{T}}(\mathbf{q})$ be r, then $\mathbb{B}_S(\mathbf{u}_1) = (a, 2r)$ and $\mathbb{B}_S(\mathbf{u}_d) = (\text{“no”}, \text{“no”})$. ☐

4.2 Knowledge Representation in Algorithms for NT^d and a Key Lemma

An algorithm for NT^d tries to learn about the connectors in $\mathcal{T}^*$ by repeatedly querying its leaves. To capture its intermediate knowledge about this $\mathcal{T}^*$, we introduce a notion of partial connectors.

Let $\sigma = [\sigma(1), ..., \sigma(k)]$ be an array of distinct elements from $\{0, 1, ...n\}$. Then, σ defines a string $S_{\sigma(1)} \circ ... \circ S_{\sigma(k)}$, referred to as a *connecting segment*. Recall $S_0 = 2, S_1 = 3 \circ 4, ...S_n = (2n+1) \circ (2n+2)$. A *partial connector* over $\mathbb{J}_n$ is then a set $\mathcal{C}$ of connecting segments such that each $j \in \mathbb{J}_n$ is contained in exactly one segment in $\mathcal{C}$ and 2 is the first element of the segment containing it. If $\mathcal{C}$ has $n+1$ segments, that is, $\mathcal{C} = \{2, 3 \circ 4, ..., (2n+1) \circ (2n+2)\}$, then $\mathcal{C}$ is called an *empty connector*. We say a connector C is *consistent* with a partial connector $\mathcal{C}$ if every segment in $\mathcal{C}$ is a substring of C.

Let $r[\mathcal{C}]$ be the last symbol of the segment in $\mathcal{C}$ that starts with 2. Let $L[\mathcal{C}]$ and $R[\mathcal{C}]$, respectively, be the set of first and the last symbols of other segments in $\mathcal{C}$. So, $r[\mathcal{C}] \in \mathbb{F}_n \cup \{2\}$, $L[\mathcal{C}] \subset \mathbb{O}_n$, and $R[\mathcal{C}] \subset \mathbb{F}_n$. Also, $|L[\mathcal{C}]| = |R[\mathcal{C}]|$. If $2 \ne r[\mathcal{C}]$, we use $\phi_{\mathcal{C}}(2)$ to denote its right neighbor. Note that each $s \in \mathbb{J}_n - L[\mathcal{C}] \cup R[\mathcal{C}] \cup \{r[\mathcal{C}], 2\}$ has two neighbors in $\mathcal{C}$. If s is even, we will use $\phi_{\mathcal{C}}(s)$ to denote its right neighbor and if s is odd, we use $\phi_{\mathcal{C}}(s)$ to denote its left neighbor.

> Initially, the knowledge of an algorithm for NT^d can be viewed as a tree $\mathcal{T}$ of empty connectors. At each round, the algorithm chooses a query point $\mathbf{q}$ and asks for $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$, which may connect some segments in the partial connectors. So $\mathcal{T}$ is updated. The algorithm succeeds when every partial connector becomes a connector and $\mathcal{T}$ grows into $\mathcal{T}^*$.

So, at intermediate steps, the knowledge of the algorithm can be expressed by a tree $\mathcal{T}$ of partial connectors.

Definition 4.10 (Valid Tree of Partial Connectors)**.** *An* (n, d)-ToPC $\mathcal{T}$ *is a tree* $T_{n,d}$ *in which each internal node* v *is associated with a partial connector* $\mathcal{C}_v$ *over* $\mathbb{J}_n$*.*

$\mathcal{T}$ *is a valid* (n, d)-ToPC *if for each internal node* $v \in T_{n,d}$ *whose children are not leaves, its partial connector* $\mathcal{C}_v$ *at* v *satisfies the following condition: For each pair* $s, t \in \mathbb{J}_n$ *with* $\phi_{\mathcal{C}_v}(s) = t$*, the tree* $\mathcal{T}_s$ *rooted at the* s^{th}*-successor* v_s *and the tree* $\mathcal{T}_t$ *rooted at the* t^{th}*-successor* v_t *of* v *are both valid ToCs, and* $\mathsf{name}\,(\mathsf{tail}\,(v_s))$ *in* $\mathcal{T}_s$ *and* $\mathsf{name}\,(\mathsf{tail}\,(v_t))$ *in* $\mathcal{T}_t$ *are the same.*

A valid (n, d)-ToC $\mathcal{T}^*$ is *consistent* with a valid (n, d)-ToPC $\mathcal{T}$, denoted by $\mathcal{T} \models \mathcal{T}^*$, if for every internal node, its connector in $\mathcal{T}^*$ is consistent with its (partial) connector in $\mathcal{T}$.

A partial connector $\mathcal{C}$ is a *β-partial connector* for $0 < \beta < 1$ if the number of segments in $\mathcal{C}$ is at least $(1 - \beta)n + 1$. To simplify our proof, we will relax our oracle $\mathbb{B}_{\mathcal{T}^*}$ to sometime provide more information to the algorithm than being asked so that the $\mathcal{T}$ it maintains always satisfies the conditions of the following definition:

Definition 4.11 (Valid (n, d, β)-ToPC)**.** *A valid* (n, d)*-ToPC* $\mathcal{T}$ *is a* valid (n, d, β)-ToPC *if its root has a* β*-partial connector, Moreover, for each internal node* $v \in T_{n,d}$ *whose children are not leaves, if the partial connector* $\mathcal{C}_v$ *at* v *is a* β*-partial connector, then it satisfies the following condition: The* s^{th}*-successor of* v*, for each* $s \in L[\mathcal{C}] \cup R[\mathcal{C}] \cup \{r[\mathcal{C}]\}$*, has a* β*-partial connector.*

Key to our analysis is Lemma 4.12 below, stating that every valid (n, d, β)-ToPC has a large number of consistent valid (n, d)-ToCs, and moreover, the names of the tails of these ToCs are nearly-uniformly distributed. Let $\mathcal{F}[\mathcal{T}] = \{\mathsf{name}\,(\mathsf{tail}\,(\mathcal{T}^*)) \mid \mathcal{T} \models \mathcal{T}^*\}$. Also, for each $\mathbf{p} \in (\mathbb{F}_n)^d$, let $N[\mathcal{T}, \mathbf{p}] = |\{\mathcal{T}^* \mid \mathcal{T} \models \mathcal{T}^* \text{ and } \mathsf{name}\,(\mathsf{tail}\,(\mathcal{T}^*)) = \mathbf{p}\}|$.

Lemma 4.12 (Key Lemma)**.** *For* $d \geq 1$ *and* $\beta \in [0, 24^{-d}]$, $|\mathcal{F}[\mathcal{T}]| \geq ((1 - \beta)n)^d$ *for each valid* (n, d, β)*-ToPC* $\mathcal{T}$*. Let* $\alpha_1(\beta) = 1$*. Then, for all* $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{F}[\mathcal{T}]$,

$$\frac{1}{\alpha_d(\beta)} \leq \frac{N[\mathcal{T}, \mathbf{p}_1]}{N[\mathcal{T}, \mathbf{p}_2]} \leq \alpha_d(\beta), \quad \text{where} \quad \alpha_d(\beta) = \frac{(\alpha_{d-1}(\beta))^7}{(2(1-\beta)^{d-1} - 1)^3}, \text{ for } d \geq 2. \tag{1}$$

Proof. When $d = 1$, let $\mathcal{C}$ be the only partial connector in $\mathcal{T}$. Clearly, $\mathcal{F}[\mathcal{T}] = R[\mathcal{C}]$. Thus, in this case the lemma is true. We will also use this case as the base of the induction below. When $d \geq 2$, let $\mathcal{C}$ be the partial connector of the root. For each $k \in \mathbb{J}_n$, let $\mathcal{T}_k$ be the subtree of the k^{th}-successor of the root. Below, we will prove by induction on d that (1) and (**) $\mathcal{F}[\mathcal{T}] = \cup_{k \in R[\mathcal{C}]}(k \circ \mathcal{F}[\mathcal{T}_k])$ are true for all d. Note that (**) and the first condition of Definition 4.11 imply that $|\mathcal{F}[\mathcal{T}]| \geq ((1 - \beta)n)^d$.

Let $\mathcal{C} = \{Y_0, Y_1, Y_2, ..., Y_m\}$ be the β-partial connector at the root of $\mathcal{T}$; assume Y_0 is the segment starting with 2. We use r_i and t_i, respectively, to denote the ending and starting symbols of Y_i. For each $k \in \{r_0, ..., r_m, t_1, ..., t_m\}$, let $\mathcal{T}_k$ denote the $(n, d-1, \beta)$-ToPC at the k^{th}-successor of the root. For each pair $(i,j) \in [0:m] \times [1:m]$ with $i \neq j$, we define

$$N_{i,j} = \textstyle\sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}_{r_i}] \cap \mathcal{F}[\mathcal{T}_{t_j}]} N[\mathcal{T}_{r_i}, \mathbf{p}] \cdot N[\mathcal{T}_{t_j}, \mathbf{p}].$$

Inductively, (1) and (**) hold for all $d' < d$. As a result, we have $|\mathcal{F}[\mathcal{T}_k]| \geq ((1-\beta)n)^{d-1}$ for every $k \in \{r_0, ...r_m, t_1, ...t_m\}$. Thus,

$$|\mathcal{F}[\mathcal{T}_{r_i}] \cap \mathcal{F}[\mathcal{T}_{t_j}]| = |\mathcal{F}[\mathcal{T}_{r_i}]| + |\mathcal{F}[\mathcal{T}_{t_j}]| - |\mathcal{F}[\mathcal{T}_{r_i}] \cup \mathcal{F}[\mathcal{T}_{t_j}]| \geq (2(1-\beta)^{d-1} - 1)n^{d-1} > 0,$$

because $\beta \leq 24^{-d}$. By the inductive hypothesis, we have $N_{i,j} > 0$, for all $(i,j) \in [0:m] \times [1:m]$ with $i \neq j$.

To show (**), it suffices to prove that $N[\mathcal{T}, \mathbf{p}] > 0$ if and only if $\mathbf{p} \in \cup_{k \in R[\mathcal{C}]}(k \circ \mathcal{F}[\mathcal{T}_k])$. Clearly, $N[\mathcal{T}, \mathbf{p}] = 0$ for $\mathbf{p} \notin \cup_{k \in R[\mathcal{C}]}(k \circ \mathcal{F}[\mathcal{T}_k])$. So, let us consider $\mathbf{p} \in \cup_{k \in R[\mathcal{C}]}(k \circ \mathcal{F}[\mathcal{T}_k])$. Since $p_1 \in R[\mathcal{C}]$, WLOG, assume $p_1 = r_m$. We use $\mathcal{P}$ to denote the set of permutations $s_0 s_1 ... s_{m-1}$ over $[0:m-1]$ with $s_0 = 0$. Then

$$N[\mathcal{T}, \mathbf{p}] = \textstyle\sum_{s_0 s_1 \ldots s_{m-1} \in \mathcal{P}} \left(\left(\prod_{i=0}^{m-2} N_{s_i, s_{i+1}} \right) \cdot N_{s_{m-1}, m} \cdot N\left[\mathcal{T}_{r_m}, (p_2, p_3, ..., p_d)\right] \right).$$

By the inductive hypothesis, every item in the summation above is positive. So $N[\mathcal{T}, \mathbf{p}] > 0$ and (**) holds for d.

Next, to prove (1), consider $\mathbf{p}_1 \in \mathcal{F}[\mathcal{T}]$ and $\mathbf{p}_2 \in \mathcal{F}[\mathcal{T}]$. There are two basic cases. When $p_{1,1} = p_{2,1}$, Eqn. (1) follows directly from (2) and the inductive hypothesis. When $p_{1,1} \neq p_{2,1}$, without loss of generality, we assume $p_{1,1} = r_m$ and $p_{2,1} = r_{m-1}$.

Let $\mathcal{P}_1$ denote the set of permutations over $\{0, 1, ..., m-2, m-1\}$ with $s_0 = 0$ and $\mathcal{P}_2$ denote the set of permutations over $\{0, 1, ..., m-2, m\}$ with $s_0 = 0$. For $P = s_0 s_1 ... s_{m-1} \in \mathcal{P}_1$, let $\Pi(P)$ be the permutation obtained from P by replacing $m-1$ by m. Clearly Π is a bijection from $\mathcal{P}_1$ to $\mathcal{P}_2$. We can write $N[\mathcal{T}, \mathbf{p}_1]$ and $N[\mathcal{T}, \mathbf{p}_2]$ as two summations:

$$N[\mathcal{T}, \mathbf{p}_1] = \textstyle\sum_{P \in \mathcal{P}_1} N_1(P), \quad \text{and} \quad N[\mathcal{T}, \mathbf{p}_2] = \textstyle\sum_{P \in \mathcal{P}_1} N_2(\Pi(P)),$$

where $N_1(P)$ and $N_2(\Pi(P))$ are given by similar terms as in (2).

We now prove for every $P \in \mathcal{P}_1$, $(N_1(P)/N_2(\Pi(P)) \leq \alpha_d(\beta)$. Let $P = s_0 s_1 ... s_{m-1}$ where $s_k = m-1$ for some $1 \leq k \leq m-1$. If $k < m-1$, then we expand $N_1(P)$ and $N_2(\Pi(P))$ as:

$$\frac{N_1(P)}{N_2(\Pi(P))} = \frac{N_{s_{k-1}, m-1} \cdot N_{m-1, s_{k+1}} \cdot N_{s_{m-1}, m} \cdot N[\mathcal{T}_{r_m}, (p_{1,2}, p_{1,3}, ..., p_{1,d})]}{N_{s_{k-1}, m} \cdot N_{m, s_{k+1}} \cdot N_{s_{m-1}, m-1} \cdot N[\mathcal{T}_{r_{m-1}}, (p_{2,2}, p_{2,3}, ..., p_{2,d})]}.$$

It then follows from the application of our inductive hypothesis to the straightforward expansion of terms $N_{i,j}$ that $N_1(P)/N_2(\Pi(P)) \leq \alpha_d(\beta)$.

Similarly, we can establish the same bound for the case when $k = m-1$. □

4.3 The Randomized Query Complexity of NT^d

By querying every leaf, one can solve any instance of NT^d with n^d queries. Below, we prove Theorem 4.4 by showing $\mathsf{RQ}^d_{\mathsf{NT}}(n) = (\Omega(n))^d$. We first relax $\mathbb{B}_{\mathcal{T}^*}$ by extending it to $(\mathbb{J}_n)^m$ for $m \in [1:d]$.

Definition 4.13 (Relaxation of $\mathbb{B}_{\mathcal{T}^*}$)**.** *Suppose $\mathcal{T}^*$ is a valid (n,d)-ToC and $\mathbf{q} \in (\mathbb{J}_n)^m$. Let v be the node with $\mathsf{name}(v) = q_1 q_2 ... q_m$. Let $\mathbf{q}' = \mathsf{name}(\mathsf{tail}(v)) \in (\mathbb{J}_n)^d$ (in tree $\mathcal{T}^*$). Then, $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q}) = \mathbb{B}_{\mathcal{T}^*}(\mathbf{q}')$.*

Query-and-Update$(\mathcal{T}, \mathbf{q})$**, where** $\mathbf{q} \in (\mathbb{J}_n)^d$

```
if $\mathcal{T}$ has complete information of $\mathbf{q}$ then return;
if $\exists\, 0 \le i \le d-1 : \left|R[\mathcal{C}_i]\right| = (1-\beta_d)n$ then
    set $m$ be the smallest of such $i$ ($m \in [0 : d-1]$)
else set $m = d$
if $m = 0$ then set $\mathcal{T} = \mathcal{T}^*$  {and $I = 1$}
else Update$(\mathcal{T}, (q_1, q_2, ..., q_m), m)$
```

Update$(\mathcal{T}, \mathbf{q}, m)$**, where** $\mathbf{q} \in (\mathbb{J}_n)^m$ **and** $1 \le m \le d$

```
fetch $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$  {set $A_m = A_m + 1$, $\mathcal{B}_m[A_m] = 0$ and $\mathcal{B}_{m,k}[A_m] = 0$}
if $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q}) = \mathcal{T}^*$ then set $\mathcal{T} = \mathcal{T}^*$  {set $\mathcal{B}_m[A_m] = 1$}
else [let $d-m \le h \le d-1$ and $r \in \mathbb{J}_n$ be the first and second components of $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$]
    set $m' = d-h-1$
    $\exists\, Y_1, Y_2 \in \mathcal{C}_{m'}$: {the ending symbol of $Y_1$, the starting symbol of $Y_2$} = $\{q_{m'+1}, r\}$
    replace $Y_1$ and $Y_2$ in $\mathcal{C}_{m'}$ by the concatenation of $Y_1$ and $Y_2$  {set $\mathcal{B}_{m,m'}[A_m] = 1$}
    let $\mathcal{T}'$ and $\mathcal{T}''$ be the third and fourth components of $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$
    replace the subtree of $\mathcal{T}$ rooted at $u_{m'+1}$ with $\mathcal{T}'$;
    replace the subtree of $\mathcal{T}$ rooted at the $r$-successor of $u_{m'}$ with $\mathcal{T}''$
```

Figure 8:

Proof (Theorem 4.4). To apply Yao's Minimax Principle [32], we consider the distribution $\mathcal{D}$ in which each valid (n,d)-ToC $\mathcal{T}^*$ is chosen with the same probability. We will prove that the expected query complexity of any deterministic algorithm $\mathcal{A}$ for NT^d over $\mathcal{D}$ is $(\Omega(n))^d$. Let $\beta_d = 24^{-d}$.

Suppose, at a particular step, the current knowledge of $\mathcal{A}$ can be expressed by a valid (n,d,β_d)-ToPC $\mathcal{T}$, which is clearly true initially, and $\mathcal{A}$ wants to query $\mathbf{q} \in (\mathbb{J}_n)^d$. Let u_0 be the root of $\mathcal{T}$ and u_i be the node with $\mathsf{name}(u_i) = q_1 ... q_i$. Let $\mathcal{C}_i$ be the partial connector at u_i in $\mathcal{T}$ and $\mathcal{T}_i$ be the subtree of $\mathcal{T}$ of u_i. There are two cases (1) $\forall i \in [0, d-1]$, $\mathcal{C}_i$ is a partial connector and $q_{i+1} \in L[\mathcal{C}_i] \cup R[\mathcal{C}_i] \cup \{r[\mathcal{C}_i]\}$. (2) otherwise. From the definition of $\mathbb{B}_{\mathcal{T}^*}$, we can show that in case (2), $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$ can be answered based on $\mathcal{T}$ only. So, WLOG, we assume $\mathcal{A}$ is smart and never asks unnecessary queries.

In case (1), because $\mathcal{T}$ is a (n,d,β_d)-ToPC, $\mathcal{C}_i$ is a β_d-partial connector for all $i \in [0, d-1]$. Let $h = \mathsf{height}(\mathsf{head}(\mathbf{q}))$. If $h = d$, then $\mathcal{A}$ gets $\mathcal{T}^*$. Otherwise, the knowledge gained by querying $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$ connects two segments in $\mathcal{C}_{d-h-1}$ and replaces the two involved subtrees by the corresponding ones in $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$. The resulting tree $\mathcal{T}$, however, may no longer be a (n, d, β_d)-ToPC, if $|R[\mathcal{C}_{d-h-1}]| = (1-\beta_d)n$ before the query. We will relax $\mathbb{B}_{\mathcal{T}^*}$ to provide $\mathcal{A}$ more information to ensure that the resulting $\mathcal{T}$ remains a valid (n, d, β_d)-ToPC. To this end, we consider two subcases: Case (1.a): if $\forall\ i \in [0 : d-1]$, $|R[\mathcal{C}_i]| > (1-\beta_d)n$, then $\mathcal{A}$ receives $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$ as it requested. Case (1.b): if $\exists\ i \in [0 : d-1]$ such that $|R[\mathcal{C}_i]| = (1-\beta_d)n$, then let $m = \min\{i : |R[\mathcal{C}_i]| = (1-\beta_d)n\}$. Let $\mathbf{q}' = (q_1, ..., q_m)$. Instead of getting $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q})$, $\mathcal{A}$ gets $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q}')$. In this way, the resulting $\mathcal{T}$ remains a valid (n, d, β_d)-ToPC. Details of the query-and-update procedure can be found in Figure 8.

We introduce some "analysis variables" to aid our analysis. These variables include: (1) $I \in \{0, 1\}$: Initially, $I = 0$. If $m = 0$ in case (1.b), then we set $I = 1$. (2) For each $m \in [1 : d]$, $A_m \in \mathbb{Z}$, and a set of binary sequences $\mathcal{B}_m[...]$ and $\mathcal{B}_{m,k}[...]$, $\forall k \in [0 : m-1]$. Initially, $A_m = 0$, and $\mathcal{B}_m$, $\mathcal{B}_{m,k}$ are empty. Each time in case (1.b) when $m > 0$, we increase A_m by 1; in case (1.a), we increase A_d by 1. To unify the discussion below, if we have case (1.a), let $m = d$ and $\mathbf{q}' = \mathbf{q}$. If $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q}') = \mathcal{T}^*$, we set $\mathcal{B}_m[A_m] = 1$ and $\mathcal{B}_{m,k}[A_m] = 0$, $\forall k \in [0 : m-1]$. Otherwise, if the first component of $\mathbb{B}_{\mathcal{T}^*}(\mathbf{q}')$ is $d - l$, for $l \in [1 : m]$, then set $\mathcal{B}_{m,l-1}[A_m] = 1$, $\mathcal{B}_m[A_m] = 0$ and $\mathcal{B}_{m,k}[A_m] = 0$ for all $0 \le k \ne l-1 \le m-1$.

Let $M_d = (\beta_d n/2)^d$. Given a random valid (n, d)-ToC $\mathcal{T}^*$, if $\mathcal{A}$ stops before making M_d queries, let $\{I, A_m, \mathcal{B}_m, \mathcal{B}_{m,k}\}$ be the set of analysis variables assigned when $\mathcal{A}$ stops; otherwise, $\{I, A_m, \mathcal{B}_m, \mathcal{B}_{m,k}\}$ is assigned after $\mathcal{A}$ makes exactly M_d queries. Let $M_i = (\beta_d n/2)^i$. We define a set of binary strings $\{\overline{\mathcal{B}}_m[1...M_m], \overline{\mathcal{B}}_{m,k}[1...M_m],\ 1 \le m \le d,\ 0 \le k \le m-1\}$ from $\mathcal{B}_m$ and $\mathcal{B}_{m,k}$: For every $1 \le i \le M_m$, (I) $\overline{\mathcal{B}}_m[i] = \mathcal{B}_m[i]$ for $i \le \min(A_m, M_m)$ and $\overline{\mathcal{B}}_m[i] = 0$ for $A_m < i \le M_m$; (II) $\overline{\mathcal{B}}_{m,k}[i] = \mathcal{B}_{m,k}[i]$ for $i \le \min(A_m, M_m)$ and $\overline{\mathcal{B}}_{m,k}[i] = 0$ $A_m < i \le M_m$.

Let $[A]$ denote that an event A is true. Let NOT-YET-FOUND $(\mathcal{T}^*)$ be the event that $\mathcal{A}$ hasn't found the tail of $\mathcal{T}^*$ after making M_d queries. Let B_m, $B_{m,k}$, $\overline{B}_m$ and $\overline{B}_{m,k}$ denote the number of 1's in $\mathcal{B}_m$, $\mathcal{B}_{m,k}$, $\overline{\mathcal{B}}_m$ and $\overline{\mathcal{B}}_{m,k}$, respectively. Then, [NOT-YET-FOUND $(\mathcal{T}^*)$] if and only if $[I = 0$ and $B_m = 0, \forall m \in [1 : d]]$. The theorem directly follows from Lemmas 4.14 below. ☐

Lemma 4.14. *Let A denote the following event,*

$$A = \left(\overline{B}_m = 0 \text{ and } \overline{B}_{m,k} \le \frac{16 \cdot M_m}{n^{m-k-1}} \text{ and } \overline{B}_{m,m-1} \le M_m, \forall\ m \in [1 : d], k \in [0, m-2] \right).$$

then **(E.1)** $[A]$ *implies* [NOT-YET-FOUND $(\mathcal{T})$] *and* **(E.2)** $\mathsf{Pr}_{\mathcal{D}}[A] \ge 1/2$.

Proof (of Lemma 4.14)*:* To prove **(E.1)**, we use the following inequalities that follow from the definition of our analysis variables.

$$1)\ \ A_m \le \frac{1}{\beta_d n} \sum_{i=m+1}^{d} B_{i,m}, \text{ for all } 1 \le m \le d-1; \quad \text{and} \quad 2)\ \ I = 1 \implies \sum_{i=1}^{d} B_{i,0} \ge \beta_d n. \qquad (2)$$

Recall that $[$NOT-YET-FOUND $(\mathcal{T})] = [I = 0 \text{ and } B_m = 0,\ \forall m \in [1 : d]]$.

To prove **(E.1)**, it suffices to show that $[A] \Rightarrow [I = 0]$ and $[A] \Rightarrow [B_m = 0,\ \forall\ m \in [1 : d]]$. We use $[A] \Rightarrow [B]$ to denote if event A is true then event B is true. It follows immediately from the definitions of B_m and $\overline{B}_m$, that if $A_m \le M_m$, then $B_m = \overline{B}_m$. So, we first inductively prove that $[A] \Rightarrow [A_{d-m} \leq M_{d-m}, \forall\, m \in [0 : d-1]]$. The base case when $m = 0$ is trivial, since A_d is at most M_d, the total number of queries.

We now consider $m \geq 1$, and assume inductively, that $A_i \leq M_i$ for all $i \in [d-m+1 : d]$. Consequently, for all $i \in [d-m+1 : d]$ and $j \in [0, i-1]$, $\overline{B}_i = B_i$ and $\overline{B}_{i,j} = B_{i,j}$. By Eqn. (2), we have

$$\begin{aligned} A_{d-m} \;\leq \sum_{i=d-m+1}^{d} \frac{B_{i,d-m}}{\beta_d n} &= \sum_{i=d-m+1}^{d} \frac{\overline{B}_{i,d-m}}{\beta_d n} \leq \frac{1}{\beta_d n}\left(M_{d-m+1} + \sum_{i=d-m+2}^{d}\left(\frac{16 \cdot M_i}{n^{i-d+m-1}}\right)\right) \\ &\leq M_{d-m}\left(\frac{1}{2} + 8\sum_{i=d-m+2}^{d}\left(\frac{\beta_d}{2}\right)^{i-d+m-1}\right) \leq M_{d-m}\left(\frac{1}{2} + 8 \cdot \frac{\beta_d}{2} \cdot 2\right) < M_{d-m}. \end{aligned}$$

Thus, $[A] \Rightarrow [B_m = 0,\ \forall\, m \in [1 : d]]$. Now we prove $[A]$ implies $[I = 0]$.

Consider the partial connector $\mathcal{C}$ at the root. We have,

$$\Big[B_m = 0,\ \forall\, m \in [1 : d] \text{ and } \textstyle\sum_{m=1}^{d} B_{m,0} < \beta_d n\Big] \Longrightarrow \big[\,|R[\mathcal{C}]| > (1-\beta_d)n\big] \Longrightarrow [I = 0]\,.$$

So it suffices to show $[A]$ implies $\big[\sum_{m=1}^{d} B_{m,0} < \beta_d n\big]$. Assume $[A]$, then

$$\sum_{i=m}^{d} B_{m,0} = \overline{B}_{1,0} + \sum_{m=2}^{d} \overline{B}_{m,0} \leq M_1 + \sum_{m=2}^{d} \frac{16 \cdot M_m}{n^{m-1}} = \beta_d n\left(\frac{1}{2} + 8\sum_{m=2}^{d}\left(\frac{\beta_d}{2}\right)^{m-1}\right) < \beta_d n.$$

The first equation follows from $[A] \Rightarrow [A_{d-m} \leq M_{d-m}, \forall m \in [0 : d-1]]$ and the first inequality uses $\overline{B}_{m,m-1} \leq M_m$ for all $m \in [1 : d]$. Finally, to prove **(E.2)**,

$$\begin{aligned} \mathsf{Pr}_{\mathcal{D}}[A] \;&=\; \mathsf{Pr}_{\mathcal{D}}\left[\overline{B}_m = 0, \overline{B}_{m,k} \leq \frac{16 \cdot M_m}{n^{m-k-1}}, \overline{B}_{m,m-1} \leq M_m, \forall\, m \in [1 : d], k \in [0 : m-2]\right] \\ &\geq\; 1 - \left(\sum_{m=1}^{d} \mathsf{Pr}_{\mathcal{D}}\left[\overline{B}_m > 0\right] + \sum_{m=1}^{d}\sum_{k=1}^{m-2} \mathsf{Pr}_{\mathcal{D}}\left[\overline{B}_{m,k} > \frac{16 \cdot M_m}{n^{m-k-1}}\right]\right) \geq \frac{1}{2}. \end{aligned}$$

The last inequality follows from Lemma 4.18. □

As $\mathcal{T}$ is chosen randomly from valid (n, d)-ToCs, $\overline{\mathcal{B}}_m$ and $\overline{\mathcal{B}}_{m,k}$ are random binary strings from a distribution defined by the deterministic algorithm $\mathcal{A}$. To assist the analysis of these random binary strings, we introduce the following definition.

Definition 4.15 (c-Biased Distributions)**.** *Suppose we have a probabilistic distribution over* $\{0,1\}^m$*. For every binary string* S *of length at most* m*, we define*

$$U_S = \Big\{ S' \in \{0,1\}^m \,\Big|\, S \text{ is a prefix of } S' \Big\}.$$

For $0 \leq c \leq 1$*, the distribution is said to be* c-biased *if we have* $\mathsf{Pr}[U_1] \leq c$ *and* $\mathsf{Pr}[U_{S \circ 1}] \leq c \cdot \mathsf{Pr}[U_S]$ *for every binary string* S *with* $1 \leq |S| \leq m-1$*.*

As an important step in our analysis, we prove the following lemma.

Lemma 4.16 (Always Biased)**.** *For all* $1 \leq m \leq d$*, the distribution over* $\overline{\mathcal{B}}_m$ *is* $2/n^m$*-biased. Similarly, for* $2 \leq m \leq d$ *and* $0 \leq k \leq m-2$*, the distribution over* $\overline{\mathcal{B}}_{m,k}$ *is* $2/n^{m-k-1}$*-biased.*

Proof. The lemma follows from Corollary 4.17 below of our Key Lemma (4.12). ☐

Corollary 4.17. *For* $d \geq 1$ *and* $\beta \in [0, 24^{-d}]$*, let* $\mathcal{T}$ *be a valid* (n, d, β)*-ToPC and let integer* $N = \sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}]} N[\mathcal{T}, \mathbf{p}]$ *be the number of consistent ToCs. For* $\mathbf{q} \in (\mathbb{J}_n)^m$ *where* $m \in [1:d]$*, if tree* $\mathcal{T}$ *has no information on* $\mathbf{q}$*, then*

1. $(N^*/N) \leq (2/n^m)$ *where* $N^* = |\{\mathcal{T}' \mid \mathbb{B}_{\mathcal{T}'}(\mathbf{q}) = \mathcal{T}', \mathcal{T} \models \mathcal{T}'\}|$*; and*
2. $(N_k/N) \leq (2/n^{m-k-1})$ *where for* $0 \leq k \leq m-2$*,* N_k *denotes the number of consistent ToCs* $\mathcal{T}'$ *such that the first component of* $\mathbb{B}_{\mathcal{T}'}(\mathbf{q})$ *is* $d-k-1$*.*

Proof. For each $k \in [0 : m-1]$, let $W_k = \{\mathbf{p} \in \mathcal{F}[\mathcal{T}_k] \subset (\mathbb{F}_n)^{d-k}$, where $p_i = q_{k+i}, \forall i \in [1 : m-k]\}$. Clearly, $|W_k| \leq n^{d-m}$. By Lemma 4.12, for all $\mathbf{p}_1$ and $\mathbf{p}_2 \in \mathcal{F}[\mathcal{T}]$, $N[\mathcal{T}, \mathbf{p}_1]/N[\mathcal{T}, \mathbf{p}_2] \leq \alpha_d(\beta)$. Thus

$$\frac{N^*}{N} = \frac{\sum_{\mathbf{p} \in W_0} N[\mathcal{T}, \mathbf{p}]}{\sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}]} N[\mathcal{T}, \mathbf{p}]} \leq \frac{\alpha_d(\beta) \cdot |W_0|}{|\mathcal{F}[\mathcal{T}]|} \leq \frac{\alpha_d(\beta) n^{d-m}}{((1-\beta)n)^d} \leq \frac{2}{n^m}.$$

The third inequality uses Proposition A.3. To prove the second statement, for $k \in [0 : m-2]$, we consider any connector C^* over $\mathbb{J}_n$ that is consistent with $\mathcal{C}_k$ and satisfies $\phi_{C^*}(Q_{k+1}) \neq$ "no". Assume $\phi_{C^*}(q_{k+1}) = r$. We use $\mathcal{T}'$ to denote the subtree of $\mathcal{T}$ rooted at the r^{th}-successor of u_k. Since $\mathcal{T}$ has no information of $\mathbf{q}$, both $\mathcal{T}_{k+1}$ and $\mathcal{T}'$ are $(n, d-k-1, \beta)$-ToPCs. Then

$$\frac{\sum_{\mathbf{p} \in W_{k+1} \cap \mathcal{F}[\mathcal{T}_{k+1}] \cap \mathcal{F}[\mathcal{T}']} N[\mathcal{T}_{k+1}, \mathbf{p}] \cdot N[\mathcal{T}', \mathbf{p}]}{\sum_{\mathbf{p} \in \mathcal{F}[\mathcal{T}_{k+1}] \cap \mathcal{F}[\mathcal{T}']} N[\mathcal{T}_{k+1}, \mathbf{p}] \cdot N[\mathcal{T}', \mathbf{p}]} \leq \frac{(\alpha_{d-k-1}(\beta))^2 \cdot n^{d-m}}{(2(1-\beta)^{d-k-1} - 1) \cdot n^{d-k-1}} \leq \frac{2}{n^{m-k-1}}.$$

☐

Lemma 4.18. *For all* $m \in [1 : d]$ *and* $k \in [0 : m-2]$*, we have*

$$\mathsf{Pr}_{\mathcal{D}}\big[\,\overline{B}_m > 0\,\big] < \frac{1}{2d^2} \quad \textit{and} \quad \mathsf{Pr}_{\mathcal{D}}\big[\,\overline{B}_{m,k} > \frac{16 \cdot M_m}{n^{m-k-1}}\,\big] < \frac{1}{2d^2}.$$

Proof. We will use the following fact: Let $\mathcal{D}^m_{IND}$ be the distribution over $\{0,1\}^m$ where each bit of the string is chosen independently and is equal to 1 with probability c. For all c-biased distribution $\mathcal{D}^m$ over $\{0,1\}^m$, for any $1 \leq k \leq m$,

$$\mathsf{Pr}_{S \leftarrow \mathcal{D}^m}\big[\,S \text{ has at least } k \text{ 1's}\,\big] \leq \mathsf{Pr}_{S \leftarrow \mathcal{D}^m_{IND}}\big[\,S \text{ has at least } k \text{ 1's}\,\big].$$

By Lemma 4.16, $\mathsf{Pr}_{\mathcal{D}}[\overline{B}_m > 0] \leq 1 - (1 - 2n^{-m})^{M_m} \leq 4(\beta_d/2)^m \leq 1/2d^2$. The second inequality uses Propositions A.2 and A.1, and the last inequality uses $\beta_d = 24^{-d}$ and the fact $m \geq 1$. We can apply the Chernoff bound [9] and Lemma 4.16 to prove the second probability bound. ☐

5 A Conjecture

We conclude this paper with the following conjecture.

Conjecture 1 (PLS to PPAD Conjecture)**.** *If* **PPAD** *is in* **P***, then* **PLS** *is in* **P***.*

6 Acknowledgments

We would like to thank Dan Spielman, Xiaoming Sun, and Xiaotie Deng for discussions that are invaluable to this work. We, especially Shang-Hua, thank Dan for strongly expressing his insightful frustration two years ago about the lack of a good measure-of-progress in equilibrium computation during our conversation about whether 2-NASH can be solved in smoothed polynomial time. We may never come up with the conjecture that led us to our main result had we not run into Xiaoming a year ago outside the entrance of the subway station by the City University of Hong Kong. Leaning against the wall between the supermarket and the subway station, Xiaoming introduced us to the work of Aldous's and Aaronson's on randomized and quantum local search. It was this conversation that started our speculation that fixed-point computation in randomized query might be harder than local search due to the lack of a measure-of-progress. We also thank Xiaoming for sharing his great insights on hiding random long paths inside $\mathbb{Z}_n^d$ in the lower bound argument for randomized and quantum local search when we were all at Tsinghua last summer. We, especially Xi, thank Xiaotie for the previous collaborations on fixed-point computation which proved to be very helpful to our technical work. We also would like to thank Stan Sclaroff for patiently teach us the punctuation rules in English when writing down a conversation.

References

[1] S. Aaronson. Lower bounds for local search by quantum arguments. In *Proceedings of the 36th annual ACM symposium on Theory of computing (STOC)*, pages 465–474, 2004.

[2] D. Aldous. Minimization algorithms and random walk on the d-cube. *Annals of Probability*, 11(2):403–413, 1983.

[3] K.J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

[4] L. Brouwer. Über abbildung von mannigfaltigkeiten. *Mathematische Annalen*, 71:97–115, 1910.

[5] X. Chen and X. Deng. On algorithms for discrete and approximate Brouwer fixed points. In *Proceedings of the 37th Annual ACM Symposium on Theory of computing (STOC)*, pages 323–330, 2005.

[6] X. Chen and X. Deng.On the complexity of 2d discrete fixed point problem. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, pages 489–500, 2006.

[7] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[8] X. Chen, X. Deng, and S.-H. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

[9] H. Chernoff. Asymptotic efficiency for tests based on the sum of observations. *Ann. Math. Stat.*, 23:493–507, 1952.

[10] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye. Leontief economies encode nonzero sum two-player games. *ECCC, TR05-055*, 2005.

[11] G.B. Dantzig. Maximization of linear function of variables subject to linear inequalities. In T.C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 339–347. 1951.

[12] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of computing (STOC)*, 2006.

[13] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag New York, Inc., 1987.

[14] K. Friedl, G. Ivanyos, M. Santha, and F. Verhoeven. On the black-box complexity of Sperner's lemma. In *15th FCT*, pages 245–257, 2005.

[15] M.D. Hirsch, C.H. Papadimitriou, and S. Vavasis. Exponential lower bounds for finding Brouwer fixed points. *Journal of Complexity*, 5:379–416, 1989.

[16] L.-S. Huang and S.-H. Teng. On the approximation and smoothed complexity of Leontief market equilibria. *ECCC, TR06-031*, 2006.

[17] T. Iimura, K. Murota, and A. Tamura. Discrete fixed point theorem reconsidered. *Journal of Mathematical Economics*, 41:1030–1036, 2005.

[18] N. Karmarkar. A new polynomial time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.

[19] V. Klee and G.J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities – III*, pages 159–175. Academic Press, 1972.

[20] C.E. Lemke and JR. J.T. Howson. Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.*, 12:413–423, 1964.

[21] R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 36–41, 2004.

[22] J. Nash. Equilibrium point in n-person games. *Porceedings of the National Academy of the USA*, 36(1):48–49, 1950.

[23] Y. Nesterov and A. Nemirovskii. *Interior Point Polynomial Algorithms in Convex Programming*, volume 13 of *Studies in Applied Mathematics*. SIAM, 1993.

[24] J.B. Orlin, A.P. Punnen, and A.S. Schulz. Approximate local search in combinatorial optimization. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 587–596, 2004.

[25] C.H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, pages 498–532, 1994.

[26] M. Santha and M. Szegedy. Quantum and classical query complexities of local search are polynomially related. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 494–501, 2004.

[27] R. Savani and B. von Stengel. Exponentially many steps for finding a nash equilibrium in a bimatrix game. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 258–267, 2004.

[28] H. Scarf. *The Computation of Economic Equilibria.* Yale University Press, 1973.

[29] E. Sperner. Neuer beweis fur die invarianz der dimensionszahl und des gebietes. *Abhandlungen aus dem Mathematischen Seminar Universitat Hamburg*, 6:265–272, 1928.

[30] D.A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of ACM*, 51(3):385–463, 2004.

[31] X. Sun and A.C-C. Yao. On the quantum query complexity of local search in two and three dimensions. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 429–438, 2006.

[32] A.C-C. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.

[33] S. Zhang. New upper and lower bounds for randomized and quantum local search. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 634–643, 2006.

A Inequalities

Proposition A.1. *For all* $\beta \geq 0$, $1-\beta \leq e^{-\beta}$.

Proposition A.2. *For all* $0 \leq \beta \leq 1/3$, $1-\beta \geq e^{-2\beta}$.

Lemma A.3. *For all* $d \geq 1$ *and* $\beta \in [0, 24^{-d}]$, $\alpha_d(\beta) \leq e^{2\cdot 24^{d-1}\beta}$.

Proof. We will use induction on d. The base case when $d = 1$ is trivial. We now consider the case when $d \geq 2$ and assuming inductively that the statement is true for all $d - 1$.

By Proposition A.2, for any $\beta \in [0, 24^{-d}]$, we have

$$\begin{aligned} \left(2\left(1-\beta\right)^{d-1}-1\right)^3 &\geq \left(2\left(e^{-2\beta}\right)^{d-1}-1\right)^3 \geq \left(2\big(1-2\beta(d-1)\big)-1\right)^3 \\ &= \Big(1-4\beta(d-1)\Big)^3 \geq \left(e^{-8\beta(d-1)}\right)^3 = e^{-24\beta(d-1)} \end{aligned}$$

By the inductive hypothesis, we have

$$\alpha_d(\beta) \leq (e^{2\cdot 24^{d-2}\beta})^7 \cdot e^{24\beta(d-1)} \leq e^{2\cdot 24^{d-1}\beta},$$

where the last inequality follows from $14 \cdot 24^{d-2} + 24(d-1) \leq 2 \cdot 24^{d-1}$, for all $d \geq 2$. $\square$